IBM

IBM Virtualization Engine, Version 1: Planning and Installation Guide

- How to use the Planning Advisor and the Information Center
- Installation examples for IBM @server platforms
- How to diagnose problems using the IBM Virtualization Engine

Christian Matthys
Hiroshi Ito
Bill Souhrada
Lynn Winkelbauer
Holger M Wittmann

Redbooks

ibm.com/redbooks

International Technical Support Organization

IBM Virtualization Engine, Version 1: Planning and Installation Guide

November 2005

SG24-6771-00

Note: Before using this information and the product it supports, read the information in “Notices” on page ix.

First Edition (November 2005)

This edition applies to the IBM Virtualization Engine Version 1.

© Copyright International Business Machines Corporation 2005. All rights reserved.
Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices ix
Trademarks x

Preface xi
The team that wrote this redbook. xii
Become a published author xiv
Comments welcome. xv

Part 1. Learn, understand, and plan before you install 1

Chapter 1. Simplifying the infrastructure 5
1.1 The complexity of today 6
1.2 Automation 7
1.2.1 Provisioning and orchestration 8
1.2.2 Autonomic computing technologies 10
1.3 Virtualization 12
1.3.1 Existing virtualization facilities. 12
1.3.2 The Web services as a new virtualization opportunity 13
1.3.3 Positioning the IBM Virtualization Engine 14
1.4 The value of the IBM Virtualization Engine 17
1.4.1 Managing infrastructure complexity 18
1.4.2 Simplifying management and reducing cost 18
1.4.3 Improving IT utilization 18
1.4.4 Improving service levels 19

Chapter 2. The IBM Virtualization Engine Systems Edition 21
2.1 Virtualization technologies for p5 and pSeries 22
2.1.1 Dynamic logical partitioning (DLPAR) and Micro-Partitioning. 22
2.1.2 IBM Cluster Systems Management. 23
2.1.3 Partition Load Manager. 23
2.1.4 Virtual I/O Server 24
2.2 Virtualization technologies for iSeries 25
2.2.1 Dynamic logical partitioning 25
2.2.2 Micro-Partitioning 26
2.2.3 Virtual Partition Manager 26
2.2.4 Virtual I/O 27
2.3 Virtualization technologies for zSeries 27
2.3.1 Virtual servers: LPAR and z/VM 27
2.3.2 Parallel Sysplex clustering 28
2.3.3 Workload Manager for z/OS 28
2.3.4 HiperSockets. 28
2.3.5 Intelligent Resource Director. 29
2.3.6 Virtual local area network 30
2.3.7 zSeries Application Assist Processor 30
2.4 Virtualization technologies for xSeries and BladeCenter 31
2.4.1 IBM Director 31
2.4.2 VMware. 31
2.4.3 Virtual Machine Manager 33

© Copyright IBM Corp. 2005. All rights reserved.

Chapter 3. The IBM Virtualization Engine Management Collection 35
3.1 IBM Enterprise Workload Manager 36
3.1.1 The EWLM topology 36
3.1.2 The domain policy 37
3.1.3 The Application Response Measurement standard 38
3.2 IBM Tivoli Provisioning Manager 40
3.2.1 The data center model 41
3.2.2 The workflow 43
3.2.3 The automation package 44
3.2.4 The deployment engine 44
3.2.5 For more information 44
3.3 IBM Director Multiplatform 44
3.3.1 The IBM Director Server 45
3.3.2 The IBM Director Agent 45
3.3.3 The IBM Director Console 45
3.3.4 For more information 46
3.4 IBM Virtualization Engine console 46
3.4.1 The Virtualization Engine console architecture components 47
3.4.2 The Virtualization Engine console user components 49
3.5 The common runtime 58
3.5.1 The components 59
3.5.2 Installing the common runtime 60
3.6 IBM Grid Toolbox V3 for Multiplatforms 61
3.7 Resource Dependency Service 62
3.8 Topology and terminology 62
3.9 Supported operating systems 63

Chapter 4. The planning advisor tool 67
4.1 Step 1: Accessing the planning advisor 68
4.2 Step 2: Completing the interview 71
4.3 Step 3: Ordering the code 73
4.3.1 Passport Advantage Online 73
4.3.2 Important notes 74
4.3.3 Summary 75
4.4 Step 4: Obtaining the code 76
4.4.1 Basic code 76
4.4.2 Fix packs 76
4.5 Step 5: Understanding your logical topology 78
4.6 Step 6: Matching the logical with the physical topology 79
4.7 Step 7: Learning about the installation process 80
4.8 Summary 81

Part 2. Installing the Virtualization Engine 83

Chapter 5. Installing IBM Director, the Virtualization Engine console, and TPM on a BladeCenter with xLinux 85
5.1 Before the installation 86
5.1.1 The topology 86
5.1.2 The installation flow 87
5.1.3 Pre-installation recommendations 87
5.1.4 Using the Virtualization Engine planning advisor 88
5.1.5 Verifying the requirements 91
5.2 The installation process 93
5.2.1 Installing the Virtualization Engine console component 93

5.2.2 Installing the IBM Director component 96
5.2.3 Installing the Tivoli Provisioning Manager component 99
5.3 Fix pack installation considerations 109
5.3.1 Applying fix packs to the Virtualization Engine console 110
5.3.2 Applying fix packs to the IBM Director 110
5.3.3 Applying fix packs to the Tivoli Provisioning Manager 110
5.4 Post-installation step: configuring the data center model 112
5.4.1 The DCM that is provided with Tivoli Provisioning Manager installation 112
5.4.2 Our DCM XML file 113

Chapter 6. Installing the Virtualization Engine console, IBM Director, TPM, and EWLM on pSeries with AIX 117
6.1 Before the installation 118
6.1.1 The topology 118
6.1.2 The installation flow 119
6.1.3 Pre-installation recommendations 120
6.2 Installing the Virtualization Engine console component 121
6.2.1 Installing and configuring Tivoli Directory Server 122
6.2.2 Installing the Virtualization Engine console code 127
6.3 Installing the IBM Director Server and Agent components 132
6.3.1 Installing DB2 Universal Database V8.1 with Fix Pack 5 133
6.3.2 Installing the IBM Director Server code on Windows 2003 133
6.3.3 Installing the IBM Director Agent code on AIX 134
6.3.4 Installing the Virtualization Engine console bridge for IBM Director 137
6.4 Installing the Enterprise Workload Manager component 139
6.4.1 Planning the installation 139
6.4.2 Installing the EWLM Domain Manager 139
6.4.3 Configuring the EWLM Domain Manager 140
6.4.4 Installing the EWLM Managed Server 143
6.4.5 Configuring the EWLM Managed Servers 144
6.4.6 Verifying the installation 144
6.5 Installing Tivoli Provisioning Manager on top of the Virtualization Engine components on a pSeries with AIX 147
6.5.1 The installation process 149
6.5.2 Step 1: Configuring system prerequisites 150
6.5.3 Step 2: Installing prerequisites utilities 150
6.5.4 Step 3: Installing and configuring the Tivoli Directory Server V5.2 code 151
6.5.5 Step 4: Installing the prerequisite software on the server node 153
6.5.6 Step 5: Applying mandatory patches 154
6.5.7 Steps 6: Installing the Tivoli Provisioning Manager code 155
6.5.8 Steps 7 to 9: Post-installation tasks 158
6.5.9 Step 10: Verifying the installation 160

Chapter 7. Installing the Virtualization Engine console, IBM Director, TPM, and EWLM on Blade with Windows 161
7.1 Before the installation 163
7.1.1 The topology 163
7.1.2 The installation flow 164
7.1.3 The pre-installation steps 167
7.2 Installing the Virtualization Engine console component 168
7.2.1 Installing and configuring the Virtualization Engine console prerequisites 169
7.2.2 Installing the Virtualization Engine console code using the wizard 180
7.3 Installing the IBM Director components 183
7.3.1 Choosing the components 183

7.3.2 Creating the IBM Director Database . . . 185
7.4 Installing Enterprise Workload Manager . . . 186
7.4.1 Configuring the EWLM Domain Manager . . . 187
7.4.2 Starting EWLM . . . 190
7.4.3 Connecting to EWLM Control Center . . . 190
7.4.4 Installing the EWLM Managed Server . . . 191
7.5 Installing Tivoli Provisioning Manager . . . 192
7.5.1 Installing the prerequisites . . . 192

Chapter 8. Installing and configuring the Virtualization Engine for IT management in an OS/400 environment . . . 201
8.1 Scenario description . . . 202
8.1.1 Objectives . . . 202
8.1.2 Topology . . . 203
8.2 Configuration steps . . . 204
8.2.1 Task 1: Completing and printing the planning worksheets . . . 204
8.2.2 Task 2: Completing the LDAP installation for the Virtualization Engine console 205
8.2.3 Task 3: Installing the Virtualization Engine . . . 211
8.2.4 Task 4: Installing and configuring the Virtualization Engine agents on managed servers . . . 219
8.2.5 Task 5: Configuring EWLM Domain Manager . . . 221
8.2.6 Task 6: Installing and configuring EWLM Managed Servers . . . 222
8.2.7 Task 7: Beginning the planning process for implementing systems provisioning 223

Chapter 9. Installing Virtual Machine Manager and VMware on managed servers . . . 225
9.1 Virtual Machine Manager . . . 226
9.2 The topology . . . 226
9.3 The VMM installation process . . . 229
9.3.1 Setting up the VMM environment . . . 229
9.3.2 Deploying an additional IBM Director Agent . . . 236
9.4 Uninstalling VMM components . . . 239
9.4.1 Uninstalling VMM on management servers . . . 239
9.4.2 Uninstalling VMM on managed servers . . . 240

Part 3. Using the Virtualization Engine Scenarios . . . 241

Chapter 10. How to use TPM and EWLM to provision a resource . . . 243
10.1 About this scenario . . . 244
10.2 The topology of this scenario . . . 245
10.3 Configuring the scenario . . . 246
10.3.1 Setting up the EWLM environment . . . 246
10.3.2 Configuring Tivoli Provisioning Manager . . . 247
10.4 Demonstration of the scenario . . . 252

Chapter 11. How to monitor an application in a heterogeneous environment using the VE console and EWLM . . . 261
11.1 The topology for this scenario . . . 262
11.1.1 Management servers . . . 262
11.1.2 Managed servers . . . 263
11.1.3 The workload . . . 264
11.2 Configuring the heterogeneous scenario . . . 265
11.2.1 EWLM consideration . . . 266
11.2.2 IBM Director consideration . . . 267
11.2.3 Virtualization Engine console consideration . . . 268

11.2.4 Tivoli Provisioning Manager consideration 274
11.2.5 WebSphere Application Server configuration 275
11.2.6 DB2 consideration 279
11.3 Validating the configuration 280
11.3.1 Validating the configuration with EWLM 280
11.3.2 Validating IBM Director and the Virtualization Engine console 280
11.4 Testing the heterogeneous scenario 281
11.4.1 Monitoring the workload running under normal conditions 281
11.4.2 Monitoring the workload under stressful condition 283
11.4.3 Resolving and monitoring the stressful conditions 286
11.5 Conclusions 288

Chapter 12. How to fix hardware constraints using TPM and IBM Director 291
12.1 The topology for this scenario 292
12.2 Configuring the IBM Director 293
12.3 Configuring RDM 293
12.4 Configuring Tivoli Provisioning Manager data center resources 294
12.4.1 Creating a resource pool 296
12.4.2 Creating the boot server 297
12.4.3 Creating the spare Blade 299
12.4.4 Creating the software stack 299
12.5 Configuring the Tivoli Provisioning Manager IBM-RDM.tcdriver and workflow 300
12.5.1 Installing or updating the IBM-RDM.tcdriver 301
12.5.2 Configuring workflow and variables on the data center resources 302
12.6 Completing the scenario: Provisioning the spare Blade resource 306

Chapter 13. How to optimize an xSeries IT environment using IBM Director and Virtual Machine Manager 309
13.1 IBM Director enhancement through VMM 310
13.2 The topology 312
13.3 Some scenarios that make use of the VMM enhancements 313
13.3.1 Topology of the physical and logical servers 313
13.3.2 VMM objects in IBM Director Console 315
13.3.3 Running the scenarios 317
13.4 Miscellaneous 338

Part 4. Summary 345

Appendix A. Planning for security 353
A.1 Planning 353
A.1.1 General consideration 353
A.1.2 Access and authentication 354
A.1.3 Data encryption 355
A.1.4 Files and application servers access (data protection) 355
A.1.5 Virtualization Engine secured networks support 356
A.2 Configuring 357
A.2.1 Virtualization Engine console 357
A.2.2 IBM Director 358
A.2.3 Enterprise Workload Manager 360

Appendix B. VMware installation 367
B.1 Setting up VMware ESX Server 367
B.1.1 Planning VMware ESX Server installation 368
B.1.2 Installing VMware ESX Server 368

B.1.3 Verifying VMware ESX Server installation 369
B.2 Setting up VMware VirtualCenter 370
B.2.1 Planning VMware VirtualCenter server installation. 370
B.2.2 Installing VMware VirtualCenter server 371
B.2.3 Verifying VMware VirtualCenter server installation. 372
B.3 Setting up VMware VirtualCenter Client 373
B.3.1 Planning VMware VirtualCenter Client installation 374
B.3.2 Installing VMware VirtualCenter Client 374
B.3.3 Verifying VMware VirtualCenter Client installation 374

Appendix C. Sample EWLM domain policy 377

Related publications 389
IBM Redbooks 389
Other publications 390
Online resources 390
How to get IBM Redbooks 391
Help from IBM 391

Index 393

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:
This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

© Copyright IBM Corp. 2005. All rights reserved.

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

@server®
@server®
Redbooks (logo)™
Redbooks (logo) ™
e-business on demand™
eServer™
iSeries™
i5/OS™
pSeries®
xSeries®
z/OS®
z/VM®
zSeries®
AIX 5L™
AIX®
BladeCenter®
CICS®
Domino®
DB2 Universal Database™
DB2®
Enterprise Storage Server®
GDPS®
HiperSockets™
IBM®
Lotus Notes®
Lotus®
Micro-Partitioning™
MVS™
NetView®
Notes®
OpenPower™
OS/390®
OS/400®
Parallel Sysplex®
Passport Advantage®
POWER™
POWER4™
POWER4+™
POWER5™
PR/SM™
Redbooks™
S/390®
Tivoli®
TotalStorage®
Virtualization Engine™
VSE/ESA™
WebSphere®

The following terms are trademarks of other companies:

EJB, Java, JDBC, JVM, J2EE, Power Management, Solaris, Sun, Sun ONE, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows server, Windows NT, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Pentium, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States, other countries, or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, and service names may be trademarks or service marks of others.

Preface

In October 2003, IBM® announced On Demand Business. On Demand Business puts the emphasis on business responsiveness and flexibility and uses Information Technology (IT) to make that happen. On Demand Business focuses on technology, not for technology's sake, but because technology enables new ways of doing business.

The On Demand Operating Environment is the computing environment that is required for building and managing an On Demand Business. It is a flexible, integrated, and reliable IT infrastructure that uniquely serves the business. The On Demand Operating Environment allows business objectives to drive the IT processes, rather than adjusting the business logic around the IT infrastructure. To become an On Demand Business, a company needs:

- Componentization of business processes to create the flexibility and innovation at the business level.
- Applications modularity to support the business processes.
- IT infrastructure simplification to be able to support the changes the business needs.

Fundamental to the IT infrastructure simplification are the concepts of automation and virtualization:

- *Automation* enables an IT infrastructure to manage many day-to-day tasks itself. With a self-managing infrastructure, efficiency is increased and resource allocation simplified. A fully automated IT infrastructure can sense changing conditions, such as surges in demand or isolated application errors, and can spot trends that could lead to costly system downtime. The infrastructure automatically responds by taking corrective actions that ensure IT resources remain aligned with business goals.
- *Virtualization* is the ability to separate the direct dependency of an application to a physical resource. Through virtualization, an enterprise:
  - Has a single, consolidated view of, and easy access to, all available resources in the network, regardless of location.
  - Efficiently accesses and manages those resources to reduce operations and systems management costs while maintaining needed capacity.
  - Responds dynamically to the application needs of its users.
  - Gathers and accesses information across the organization quickly to gain competitive advantage.

The IBM Virtualization Engine™ was announced and made available in August 2004 and addresses the virtualization aspects of the On Demand infrastructure.

This IBM Redbook provides:

- Information on how to prepare for the installation of the IBM Virtualization Engine Version 1 in Part 1, "Learn, understand, and plan before you install" on page 1.
- Examples of how to install the IBM Virtualization Engine Version 1 (an example is available for each type of IBM @server®) in Part 2, "Installing the Virtualization Engine" on page 83.
- Examples of how to use the IBM Virtualization Engine Version 1 to fix customer pain points in Part 3, "Using the Virtualization Engine Scenarios" on page 241.
- Hints and tips from our installation experiences in Part 4, "Summary" on page 345.

© Copyright IBM Corp. 2005. All rights reserved.

The team that wrote this redbook

This redbook was produced by a team of specialists from around the world working at the International Technical Support Organization (ITSO), Poughkeepsie Center.

Christian Matthys spent more than 25 years at IBM as a System Engineer, working with large mainframe-oriented customers. He spent three years as an ITSO project leader on assignment in Poughkeepsie, NY. In 2000 he joined the EMEA Design Center for On Demand Business, working with customers' projects to make use of the leading edge technologies, in particular the Autonomic Computing technologies. He was part of the launching team of the Start Your Engines workshops to promote the IBM Virtualization Engine. He works as an ITSO project leader based in the Product and Solutions Support Center, based in Montpellier, France. He is a Consulting IBM IT Specialist certified.

Hiroshi Ito is an IT Specialist from IBM Japan Systems Engineering. Hiro works for Virtualization Technologies such as Grid Computing for 2 years. As a Grid Computing specialist, he participated in different proofs of concept projects and customer projects. Before being in charge of Virtualization Technologies, he was an AIX® technical support in Japan. He is a Red Hat Certified Engineer.

Bill Souhrada is an IT Specialist with IBM Global Services Server Systems Operations (SSO). He has been an Intel® server Team Leader for IBM Strategic Outsourcing accounts with a primary focus on Microsoft® Windows® Server solutions. He is a Microsoft Certified Professional, and has over 9 years experience architecting, and implementing Windows Server and other IT technologies. Bill holds a BS in Health Systems Management from the University of Connecticut. He is based in Southbury, CT.

Lynn Winkelbauer is an IT Specialist for IBM Poughkeepsie. She has 21 years of experience in the zSeries® area, working in Performance, Parallel Sysplex®, and e-business on demand™. She holds a Bachelor of Science degree in Computer Science and Business Administration and is an IBM Certified Specialist in the IBM @server zSeries Technical area. Her areas of expertise include Linux®, CICS® Transaction Server, WebSphere®, performance analysis and the IBM Virtualization Engine.

Holger M Wittmann is a certified IT Architect, based in Stuttgart, Germany. He started working for IBM 1991 in the scientific center in Heidelberg, Germany, building mobile radio software solutions. In 1998 he joined IBM Global Service to build in various customer projects Java™ and Lotus® Notes® solutions. Three years later he took an opportunity in Sales & Distribution to support customers and major Independent Software Vendors (ISVs) on an architectural level in the arena of e-Procurement and Supply-Chain-Management solutions. Three years ago he joined the IBM Grid team to work with new customers implementing grid solutions and as a logical evolution to grid the theme virtualization came to his portfolio. Furthermore, he is a teacher and developer of courseware for the technical community on the topic On Demand Business with an emphasis on business architecture comprising business transformation aspects with IT Architecture.

Figure 1 The team who wrote this book, from left to right: Bill, Hiro, Lynn, Holger, and Christian

A very special thanks to the following people who made this project happen:

Madeline Nick
IBM Systems Group Technical Support

Rich Conway and Octavian Lascu
International Technical Support Organization, Poughkeepsie Center

Thanks to the following people for their contributions to this project:

Philippe Fernandes
IBM France, who did the security tests and wrote the Appendix A, “Planning for security” on page 353

Dave N Amundson
IBM US

Paola Bari
IBM US

Korin J Bevis
IBM US

James J Bozek
IBM US

Mike Connolly
IBM US

John Hock
IBM US

Tim O’Keefe
IBM US

Emmanuel Lieurain
IBM France

Thomas Monza
IBM US

Paul Oldfield
IBM US

Lydia Parziale
IBM US

Jeffrey Prentice
IBM US

John C Rossow
IBM US

Gregory Sansoni
IBM US

Richard Sit
IBM US

Brad Smoley
IBM US

Deva Walksfar
IBM US

Joe Winkelbauer
IBM US

Become a published author

Join us for a two- to six-week residency program! Help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners, or customers.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you'll develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our Redbooks™ to be as helpful as possible. Send us your comments about this or other Redbooks in one of the following ways:

- Use the online **Contact us** review redbook form found at:

  ibm.com/redbooks

- Send your comments in an email to:

  redbook@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. HYJ; HYJ Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

Part 1

Learn, understand, and plan before you install

Before you implement the IBM Virtualization Engine in your infrastructure, there are important planning, ordering, and deployment considerations. The IBM Virtualization Engine can be built of multiple components. Each component involves multiple subcomponents, and in some cases, multiple operating systems might be required.

The IBM Virtualization Engine does not require complex installation tasks. Many processes are already integrated, and more common subcomponents will be integrated with the upcoming new release of IBM Virtualization Engine. However, the functions that are put in place by the IBM Virtualization Engine are not trivial. The concepts are easy to phrase (such as provisioning, end-to-end monitoring, management of pools of resources, and so forth) and more difficult to set up in an existing environment. Understand that the complexity comes from the number of already installed standards, servers, habits and so forth, not from the tool itself.

The IBM Virtualization Engine aims to manage an heterogeneous environment. A good use of the multiple functions might require the participation of multiple operating systems skills, which need to work together in a common project.

Good planning involves identifying the problem that you want to solve and having a good understanding of the components of the solution that you want to deploy. The IBM Virtualization Engine initiative proposes different activities and tools to help you plan the best solution for you. Some of them are summarized in Figure 1 on page 2. You do not need to go

© Copyright IBM Corp. 2005. All rights reserved.

to all of them and you can choose from the list the most appropriate depending on your own environment and knowledge.

Note: The available activities and tools can vary, depending on time and countries. Contact your IBM representative to get the latest local available facilities.

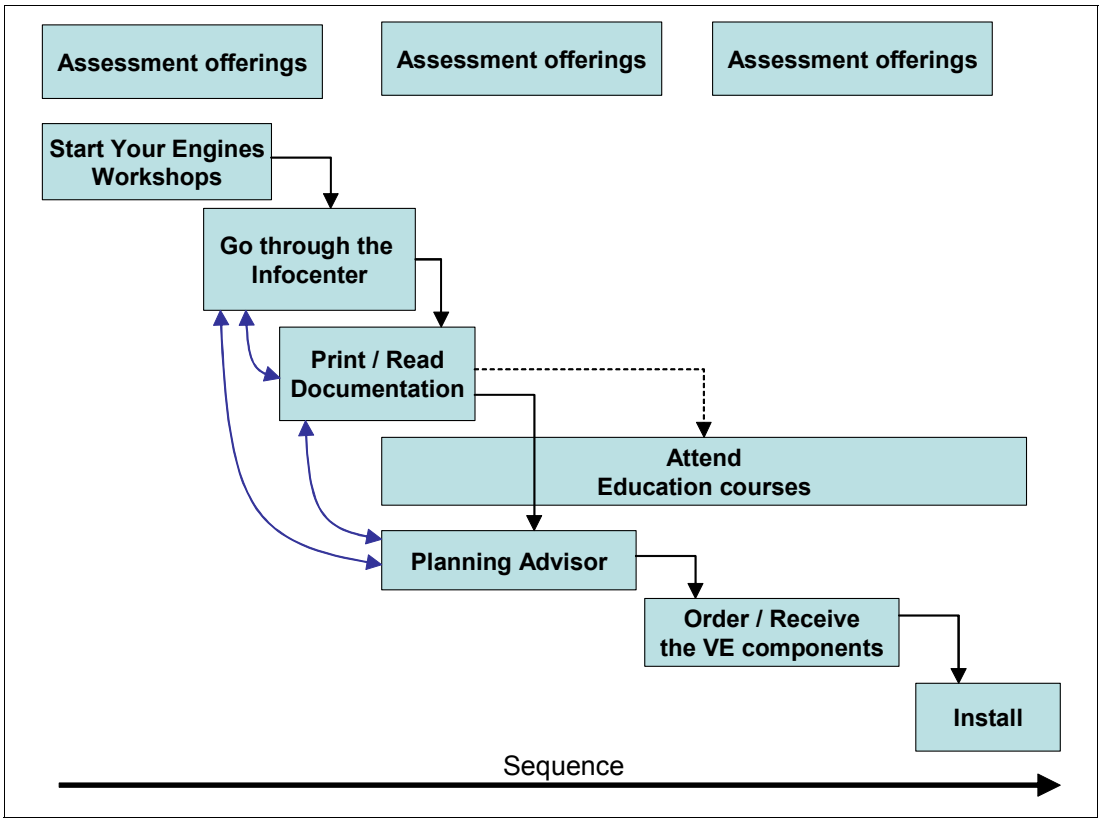

Figure 1 Planning for the installation of the IBM Virtualization Engine

As you work through the planning process, you should consider the following:

- Multiple consulting offerings are available from IBM, from very short *assessments*, where virtualization possible solutions are identified to fix specific pain-points, to more complete consulting assessments to build a virtualization road map at the enterprise level, including the evaluation of business benefits. For more information ask your IBM representative.
- The Start Your Engines (SYE) workshops provide an introduction to the IBM Virtualization Engine, providing hands-on lab exercises from which the participants can experience the components directly. For the U.S., you can refer to the following Web site and then search on course code VE001 and VE011, or ask your IBM representative in other countries:

  `http://www.ibm.com/training/us`
- The IBM @server Software Information Center is another starting point to get the latest information and gives you access to the latest documentation. The center is available at:

  `http://www.ibm.com/servers/library/infocenter`

  The SYE workshops are not a requirement before you look at the Information Center and you can start reading right away. However, the SYE workshops will give you some practice.

The IBM @server Software Information Center is named the *Information Center* in the remainder of this book.

- Though obvious, we recommend that you go through the documentation and learn about the multiple topologies, reading either from the Information Center or from printed documentation.
- Education courses are also available to go into depth for each of the components. Refer to the local IBM education Web site to get the latest information.
- When you understand better the concepts, the solutions, the components of the solutions, and their integration, the planning advisor can help you build the topology based on your requirements, can guide you in planning the installation, and can help you in the installation process.

This part of the redbook:

- Introduces the role that virtualization plays in simplifying the IT environment, in Chapter 1, “Simplifying the infrastructure” on page 5.
- Describes in details the IBM Virtualization Engine components, in Chapter 2, “The IBM Virtualization Engine Systems Edition” on page 21 and Chapter 3, “The IBM Virtualization Engine Management Collection” on page 35.
- Explains how to use the Planning Advisor, a wizard that helps configure the IBM Virtualization Engine in your environment, in Chapter 4, “The planning advisor tool” on page 67.

1

Simplifying the infrastructure

To support business priorities, the IT infrastructure needs to have the speed, the flexibility, and the resilience to handle a broad spectrum of concurrent processes to take action according to the business needs. This chapter addresses the challenges of today's IT infrastructure.

© Copyright IBM Corp. 2005. All rights reserved.

1.1 The complexity of today

Complexity is likely the most significant barrier to building the IT environment to support business growth objectives. Complexity did not appear overnight. It was an evolutionary process started years ago. More servers, more storage, and more applications were added to support the business needs, resulting in a variety of operating platforms in specific standards in multiple physical locations.

Proprietary technologies and lack of open standards obstructed integration, interoperability and unified management. Dedicated resources cannot be shared among business processes because of the barrier of the heterogeneity. Without a unified management view of the resources, IT skills must focus on IT issues instead of business priorities. Furthermore the existence of these disparate systems triggers new problems, such as security, accessibility, and performance, that in turn require new IT investments to fix IT problems.

Complexity has a price and can have negative business consequences, such as the following:

- Systems might be under utilized because they do not take advantages of performance capabilities.
- The IT infrastructure can include unused resources that are bought in case of problems, or the infrastructure might be unable to connect to other resources that face the lack of standards cooperation.
- The IT infrastructure can include resources that are used fully, meaning that the infrastructure cannot handle properly spikes in usage.
- Complexity can affect quality of service, availability, response time, and so forth.

The IBM On Demand Operating Environment strategy centers around automation, integration, and virtualization to enable a more competitive business that uses advanced technology to simplify the environment. The On Demand Operating Environment architecture, illustrated in Figure 1-1 on page 7, is based on the following fundamental capabilities:

- *Integration*, which means that everything connects through a service-oriented architecture. A service-oriented architecture (SOA) is an application framework that takes everyday business applications and breaks them down into individual business functions and processes, called services. An SOA lets you build, deploy, and integrate these services independent of applications and the computing platforms on which they run.
- *Infrastructure management*, which lets you simplify and optimize complex IT environments. Infrastructure services create the environment where business services are deployed. They also manage the infrastructure. Unlike integration services, infrastructure management is built by an on demand infrastructure and application builders. Infrastructure services include utility business services as well as service level automation and orchestration. They also include resource virtualization services.

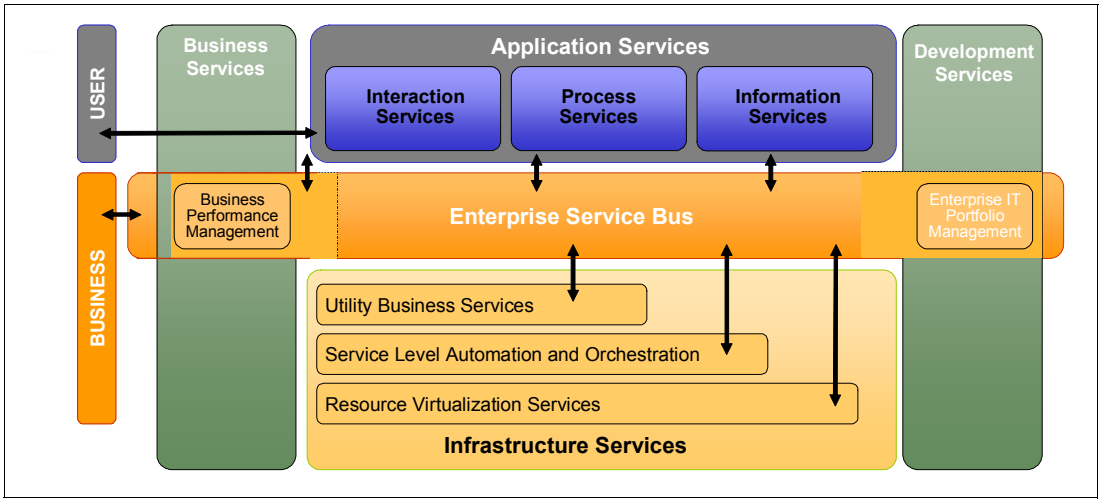

Figure 1-1 The On Demand Operating Environment architecture framework

To counter the complexity, you can use simplification techniques, such as server consolidation, remote management, and others already. Automation and virtualization are the two key elements to simplifying an IT infrastructure environment.

1.2 Automation

Infrastructure automation reduces the complexity of management, enables better use of assets, and improves availability and resiliency. Automation consists of the following elements:

- *Availability* helps ensure the health and appropriate functioning of IT environment. Availability does the following :
  - Increases infrastructure resilience.
  - Reduces time for problem identification and resolution.
- *Security* helps ensure that information assets are confidential and that data integrity is protected. Security does the following:
  - Provides enhanced protection from unauthorized access by identifying and responding automatically to security threats.
  - Promotes regulatory compliance by protecting and retaining data.
  - Saves time by providing central administration of user rights and policies.
- *Optimization* helps ensure the most productive utilization of IT resources. Optimization does the following:
  - Increases storage resource use by enabling storage capacity management.
  - Increases IT resource use by creating resilient workload management infrastructures.
  - Increases efficiency and responsiveness of batch operations by centralizing control of interoperating jobs.

- *Provisioning* makes available the right resources to the right processes and people by tracking use by application, by setting thresholds, and by executing workflows automatically to provision additional resources as needed. Provisioning does the following:
  - Reduces resource over-capacity.
  - Saves time by building, deploying, and reclaiming servers and associated network, storage, and software automatically.
  - Reduces errors.
  - Improves productivity.

These elements are linked together through *Policy-Based Orchestration*, which provides a high level of automation to coordinate IT resources to support the implementation of business policies across the infrastructure. In order to automate the management of the environment to meet business goals successfully, these service-level objectives must be defined and available to the automation component of the framework. The Business Services Management component is where these goals are defined.

The IBM Tivoli® family of products provides many solutions regarding automation. Many documents are available that provide technical details about automation solutions. The following Web sites are a good starting point to find the documentation that meets your needs:

```
http://www-18.lotus.com/wps/portal/automation
http://www-306.ibm.com/software/tivoli/features/may2004/index.html
http://publib-b.boulder.ibm.com/Redbooks.nsf/Portals/Tivoli
http://www.redbooks.ibm.com
```

The following major aspects of the automation portfolio are of particular interest in this redbook:

- Provisioning and orchestration

  The virtualization of resources gets additional value when these resources can be used dynamically and intelligently.

- Autonomic computing technologies

  To be efficient, some virtualization tasks need to be managed by specific engines without it being obvious to the user.

1.2.1 Provisioning and orchestration

Many IBM technologies, functions, offerings, tools, and products address provisioning and orchestration. These solutions are available at the server level with no possible interaction among different types of servers:

- IBM Tivoli Provisioning Manager
- IBM Director Remote Deployment Manager (RDM)
- AIX Network Installation Manager (NIM)
- Cluster System Management (CSM)
- IBM @server - Capacity Upgrade On Demand (CUoD)
- TotalStorage® Productivity Center
- Tivoli Storage Area Network Manager
- IBM Tivoli Intelligent Orchestrator
- Web Infrastructure Orchestration
- WebSphere Extended Deployment (WebSphere XD)
- Server Allocation for WebSphere Application Server
- Deep Computing Capacity on Demand (DCCoD)

Provisioning and orchestration are complementary concepts and are different solutions or products. In this redbook, we present only two key IBM solutions which address an heterogeneous environment and sometimes make use of a specific server or a specific technology:

- IBM Tivoli Provisioning Manager
- IBM Tivoli Intelligent Orchestrator

Tivoli Provisioning Manager

Tivoli Provisioning Manager, through IT Service Management workflows, consistently automates the manual provisioning and deployment process handling automatically with errors. It uses prebuilt *industry best practice* workflows to provide control and configuration, as well as allocation and reallocation of major vendors' products. User-customized workflows can implement a company's data center best practices and procedures. IT service management workflows consist of the automated steps that are executed dynamically to carry out a provisioning operation. Workflows capture and automate best practice processes for monitoring, sensing, and responding to demands for computing capacity based on business requirements. Examples of workflows include:

- Allocating and removing servers to and from clusters
- Installing and uninstalling software and patches
- Configuring servers and network devices
- Provisioning storage
- Performing bare-metal builds
- Testing floor automation

The IBM provisioning solution includes prebuilt workflows. It provides tools for customers and IBM business partners to incorporate easily into their own best practices. It does not provide automatic monitoring, analyzing, or planning functionality.

For more information about how to use and develop workflows, refer to *Developing Workflows and Automation Packages for IBM Tivoli Intelligent ThinkDynamic Orchestrator*, SG24-6057.

Workflows can be shared among users through the Orchestration and Provisioning Automation Library (OPAL). From OPAL, you can get predefined developed workflows that are dynamic and that can grow over time. You can find more information about the OPAL initiative and how to get workflows at:

http://www.developer.ibm.com/isv/tivoli/workflow.html

Note: Systems provisioning is one of the IBM Virtualization Engine systems services Version 1 as described in 3.2, “IBM Tivoli Provisioning Manager” on page 40. Systems provisioning is delivered by Tivoli Provisioning Manager 2.1.

Tivoli Intelligent Orchestrator

Tivoli Intelligent Orchestrator monitors the servers, middleware, and applications under its control, senses degrading performance, and determines actions to take. It can determine where (for which application) a resource is needed and can instruct the Tivoli Provisioning Manager to deploy a server automatically, install the necessary software, and configure the network (and deprovision when not needed any more, making the resource available for another process). Using its capacity management capabilities, Tivoli Intelligent Orchestrator can predict resource availability or need and begin the provisioning process, on demand, to help match IT resources with fluctuating workloads.

Tivoli Intelligent Orchestrator extends the benefits of the Tivoli Provisioning Manager, issuing instructions dynamically to Tivoli Provisioning Manager, which then uses automation

packages to maintain server availability and to meet required service levels in accordance with business priorities.

Note: Tivoli Intelligent Orchestrator is *not* part of the Virtualization Engine systems services.

1.2.2 Autonomic computing technologies

The highest level of automation is achieved through *autonomic computing technologies*. Autonomic computing technologies are intelligent resources that know themselves, continuously tune themselves, adapt to unpredictable conditions, prevent and recover from failures, and provide a safe environment. Autonomic computing provides the technology to enable information systems to be self-managing. These self-managing characteristics combine to deliver the automation required of an On Demand Operating Environment.

To automate an heterogeneous environment, the components need to understand each other. To understand each other, they need to speak a common language. Thus, the autonomic computing IBM initiative focuses on common formats and common standards. The autonomic computing technologies provide multiple core capabilities that can be embraced by all products. Autonomic computing technologies provide the following:

- *Data collection* provides a common logging and tracing format — the Common Base Event (CBE) format — and addresses the issue of the complexity of problem determination in multi-component systems. It introduces standard interfaces and formats to correlate system status or problem information from multiple data sources.
- *Solution installation* provides a common install technology across all our products. This common install technology will reduce deployment time as well as software maintenance time. The goal is to have a *solution-ready* install that understands and resolves dependencies to ensure the successful installation of an entire solution.
- *Policy and Information Model* ensures that each element of the system described through specific policies is acting in concert with the whole with the context of business policies. IBM is developing a common set of specifications for expressing and communicating policies as well as common technologies for entering, storing, and updating policies.
- *Resource Provisioning* provides the ability to add or remove resources capacity, such as computing power, storage, and so forth. Provisioning additional resources requires an orchestrator as well as technologies in each product that can respond to provisioning requests.
- *Common console* provides a common format and common repository for providing information about the operation of the system. An Integrated Solution Console is one of these. It is a window into their management that is being developed for each product, with one consistent user interface, a common runtime infrastructure, and a set of development tools based on industry standards and component reuse.

Note: In particular, as described in “The Integrated Solutions Console” on page 48, the Virtualization Engine console is based on this common console format.

These technologies are not fully completed yet. They will mature and will be integrated more and more into products in the future. Using standards, they are the key technologies to make the components work seamlessly.

For more information about how to start with autonomic computing technologies, refer to *A Practical Guide to the IBM Autonomic Computing Toolkit*, SG24-6635, or go to the following:

http://www-03.ibm.com/autonomic/index.shtml

The MAPE loop

The IBM autonomic computing technologies are based on a loop architecture approach. This retro-feedback loop describes all the needed steps for an infrastructure management product to apply changes into the environments. This loop is commonly called the *MAPE loop*, as an acronym for the names of the four steps that are used: Monitor, Analyze, Plan, and Execute, as shown in Figure 1-2.

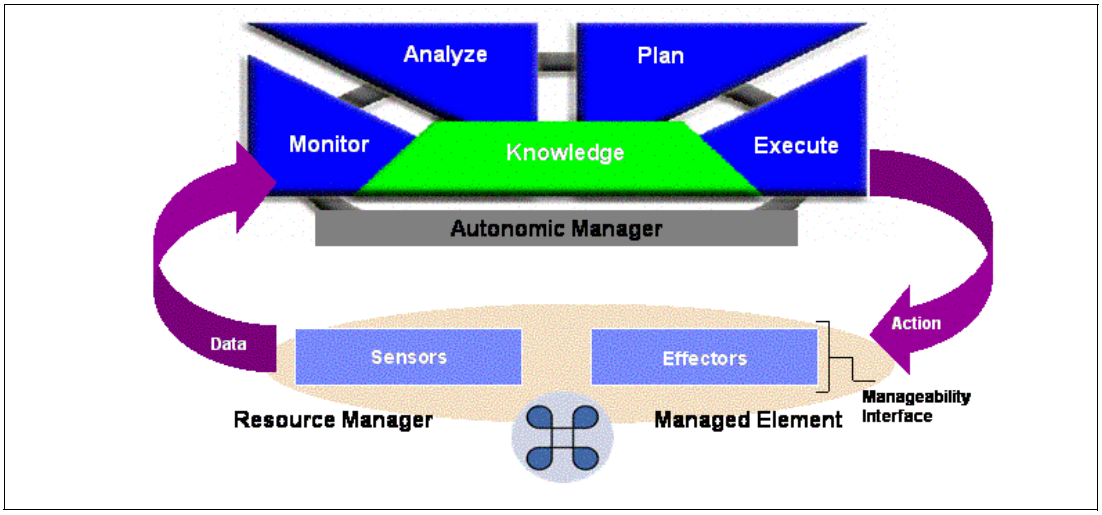

Figure 1-2 The MAPE loop

The steps of the MAPE loop are as follows:

1. The *monitor* entry uses sensors to collect information (for example the response time, a consumption number, and so forth).
2. This information needs to be *analyzed* to generate a result that is compared with a predefined objective before deciding if this information is good or needs an action to make it better (for example the response time is higher than expected).
3. When the objective is not met, a solution needs to be *planned* to make this response time better. Different solutions can be tested, in theory, that estimate the effects of a particular solution on other components. Then the solution has to be chosen. One such example might be: I will add a new CPU to give more resources; that requires to add a new server because all the cps in the running boxes are already busy.
4. The solution needs to be implemented through effectors using the *execute* step (for example from the free pool resources in front of me, I power on this server on which I need to install an operating system and the middleware needed to run the application).

A knowledge repository allows all the components of the MAPE loop to share and build data.

The IBM virtualization components as described in the sections that follow maps this MAPE loop architecture.

1.3 Virtualization

The automation concepts described previously become more effective when some virtualization capabilities are implemented. The virtualization architecture builds the base layer for the On Demand Operating Environment architecture as shown in Figure 1-1 on page 7.

Virtualization enables basic systems management of multiple, often heterogeneous systems right out of the box. It deploys and optimizes IT resources in real-time and takes a solid open standards-based approach to connectivity offering a single dashboard to pull it together. Virtualization involves a shift in thinking from physical to logical, treating IT resources as a logical pool rather than as separate physical entities. Virtualization consolidates and pools IT resources and provides a single system illusion for servers, storage, distributed systems, and networks. With virtualization you can do the following:

- Manage multiple, unlike systems.
- Deploy and optimize IT resources in real time dynamically.
- Operate the infrastructure on an open standards-based approach.

Virtualization can give new capabilities that will change:

- How resources are viewed.
- How resources are allocated.
- How data are managed aligning cost with business value.

1.3.1 Existing virtualization facilities

In most cases, virtualization technologies are already incorporated into existing infrastructure via mails servers, Web applications, storage area networks, which most often consists of pools of resources. Virtualization provides a logical rather than a physical view of resources: server, network, and data. Various and many technologies can be used to help businesses better leverage their servers, networks, and storage infrastructures.

For example, IT departments can partition some servers to increase their utilization through logical partitioning capabilities (as they exist today for the zSeries, the pSeries®, and the iSeries™). Other options include using Blade technology to enable better management of Intel-based systems, taking advantage of networking capabilities that are built into servers today (such as HiperSockets™ and Virtual LANs), or exploiting the capabilities of current storage technologies such as block virtualization, file aggregation, and centralized management.

In the following sections, we describe a few of them as examples of key facilities for linking virtualization and automation capabilities.

Logical partitioning

Logical partitioning (LPAR) is a feature that allows multiple instances of operating systems to run concurrently on the same physical hardware. Each LPAR has its own installed instance of the operating system. If there is a problem that affects the operating system, the problem effects only the software and applications installed on that one affected partition.

The zSeries mainframes introduced the concept of LPARs in the 1980s. Today, logical partitions are also supported on the pSeries and iSeries.

LPARs have many benefits in a virtualized environment. For example, with LPARs, you can isolate applications and build a test environment where you can test new applications without risk to the key business applications. You can also give resources to the most critical

applications because the LPARs resources (CPU, memory, or I/O) can be redistributed dynamically among all the LPARs.

The Virtual Partition Manager (VPM) introduces new facilities on iSeries with the recently announced POWER5™ processor.

Virtual machines

Virtual machines are similar to logical partitions in that they create a high degree of isolation between workloads running on the same physical machine. The virtual machine operating system (z/VM®) on the IBM zSeries mainframe has provided this option for decades, offering the possibility to define hundreds or thousands of virtual machines on one single server. VMware is a virtual machine technology that supports the xSeries® platform.

Virtual machine emulators are more flexible but can be less resilient than the LPARs, because the overall stability of the system is limited to that of the host operating system. The flexibility comes with the number of instances and ease of management of the virtual machine. The number of CPUs does not limit the number of instances that can be supported. For example, an Intel 2-way machine can have up to 16 virtual machines. This type of configuration is ideal for test environments where transaction volume testing is not a concern.

Workload manager

Inside a *box*, workloads can negatively affect the performance of other workloads. To prevent resource starvation of important workloads, workload management techniques can be useful for large SMP servers or even uniprocessor machines and can be partitioned. Workload management provides the control to ensure system resources are provided to the application that are most critical to the business. Workload management facilities exist for the zSeries, the pSeries and the iSeries. Each server has its own solution (for example, workload management in z/OS® and Partition Load Manager (PLM) in AIX). A workload manager solution needs to be implemented across platforms to address the heterogeneous environment of an IT infrastructure.

Enterprise Workload Manager (EWLM), new in the Virtualization Engine Version 1 and the Automated Response Measurement (ARM) standard introduce new facilities in the workload management area. These new facilities are described in 3.1, “IBM Enterprise Workload Manager” on page 36.

1.3.2 The Web services as a new virtualization opportunity

Web services initially focused on Web applications (even small sub-components of applications) to externalize pieces of existing applications and offer them to other applications or users to build new applications. Web services has now expanded its scope to include a new range of possibilities for how organizations and their partners can cooperate and connect their business applications across a network (inside and outside the enterprise), regardless of platform, design, or runtime environment.

Web services have demonstrated their capabilities to tie dissimilar resources into a stronger unified platform, using common formats and a common language. Web services provide a thin, universal layer for common communications using a consistent and standardized interface allowing for integration across platforms, operating systems, and applications.

Fundamentally, all of the capabilities of the on demand infrastructure management rely on the ability to discover, compose, and interact with the resources and the managers which control them. Many types of resources exist in the on demand operating infrastructure, including hardware (such as processors, network devices, and storage), operating systems, middleware, and applications.

The on demand environment must provide the capability to dynamically create new resources, based on business context, by using existing resources from the environment. For example, the environment must enable the creation of additional logical partitions within a server and must drive the software installation and configuration that is required to put the new application into service.

At the most basic layer, the On Demand Operating Environment must provide resources that have interfaces with managers that control the state and the function of the resource. The On Demand Operating Environment must provide standard connectors to existing native instrumentation technologies and products.

The Web services are now extending, becoming the fabric of how we build, connect, and integrate not only applications but servers and devices. Web services are used to discover computer systems, to determine their basic characteristics, to determine their state and relationship to other computer systems, to subscribe to state change notifications, and to perform basic control operations.

Standardization provides the basis for inter operability between components by codifying interfaces between components. Standardization provides the opportunity for different vendors to develop implementations that exploit the standard platform interfaces. The following are standards of particular interest:

- The OASIS[1] Web Services Distributed Management (WSDM) is responsible for defining how to represent the manageability interfaces of IT resources as Web services interfaces. The WSDM specifications depend on other Web services foundation specifications being standardized in OASIS such as WS Resource Framework specifications (WS-RF).
- DMTF[2] defines and develops the Common Information Model (CIM). CIM's focus is on providing end-to-end consistent management abstractions to be used by monitoring and traditional management systems.

To understand better the motivation between the development of these standards, refer to the following:

```
http://www-128.ibm.com/developerworks/webservices/library/ws-wisdom/
http://design.ibm.com/software/info/topic/webservices.html
http://www.oasis-open.org/home/index.php
http://www.dmtf.org/home
```

1.3.3 Positioning the IBM Virtualization Engine

To avoid the *big bang* approach that would require the re-implementation of all existing resources in the system, an incremental and staged approach to standardize the management behavior of resources in the On Demand Operating Environment is required. Existing resource implementations must be extended and adapted to allow the integration with the on demand management infrastructure in compliance with the proposed standards.

Available since August 2004, the IBM Virtualization Engine is a comprehensive portfolio of systems technologies and tools that can help you aggregate existing pools of resources and get a consolidated view of them throughout your IT environment. It uses key IBM virtualization technologies to give you a logical — rather than physical — view of data, computing power, storage capacity, and other resources. The IBM Virtualization Engine can help you automate the management of these resources based on business goals and make basic systems management of multiple disparate systems possible.

[1] The Organization for the Advancement of Structured Information Standards (OASIS) is a not-for-profit, global consortium that drives the development, convergence and adoption of e-business standards.
[2] The Distributed Management Task Force, Inc. (DMTF) is an industry organization working for the development of management standards and integration technology for enterprise and Internet environments.

The IBM Virtualization Engine wants to bridge the existing and very valuable virtualization technologies to new emerging standards and more automation processes, making the virtualization techniques a powerful IT asset for the business benefits. With the IBM Virtualization Engine, you can start to put in place the technology that enables real-time, dynamic deployment and optimization of IT resources. Its open standards-based approach allows you to manage all your systems from a single, easy-to-use dashboard.

Implementing virtualization solutions does not mean the entire IT infrastructure has to be changed in a single large-scale re-engineering project. IBM offerings that are available today are as follows:

- The IBM Virtualization Engine for Systems, which focuses on using the technologies that are integrated into IBM @server (such as Hypervisor, VLAN, Virtual I/O, and so forth).
- The IBM Virtualization Engine for Enterprise, which provides a collection of multiplatform system tools to help to simplify the management and utilization of IT resources (such as Enterprise Workload Manager, Tivoli Provisioning Manager, IBM Director Multiplatform, Virtualization Engine console, and the IBM Grid Toolbox V3 for Multiplatforms).

The evolution of virtualization

In 2005, IBM has taken steps in virtualization by transforming the Virtualization Engine into an integrated *platform* of technologies that broadens and enhances the current capabilities. The traditional view tends to focus on partitioning. This new view expands virtualization to allow multiple systems to perform together in a larger context and to give a common way to address component-level devices. Components of this system include the following:

- *Virtual resources* include all IBM server and storage systems (as well as selected servers and storage systems from other vendors such as HP, Sun™, and EMC). They use virtualization technologies that are integrated and delivered within the IBM servers (such as Hypervisor, Virtual I/O, VLAN, and so forth). These embedded technologies allow resource virtualizers to facilitate the virtualization of resources within a server or storage device.
- The *Virtualization Engine management* provides workload, performance and resource managers as well as modeling, mapping, and dependency capabilities. It includes Enterprise Workload Manager, IBM Director, and Resource Dependency Services.
- The *Virtualization Engine virtual access* point includes a virtual view and a programmatic interface. The virtual view includes the Virtualization Engine console, a single portal for administering all virtual resources such as resource health, monitoring, and integration access to planning tools. The programmatic interface offers access into the Service Oriented Architecture, grid services, and resource modeling.

Figure 1-3 on page 16 is a schematic view of the updated Virtualization Engine offering.

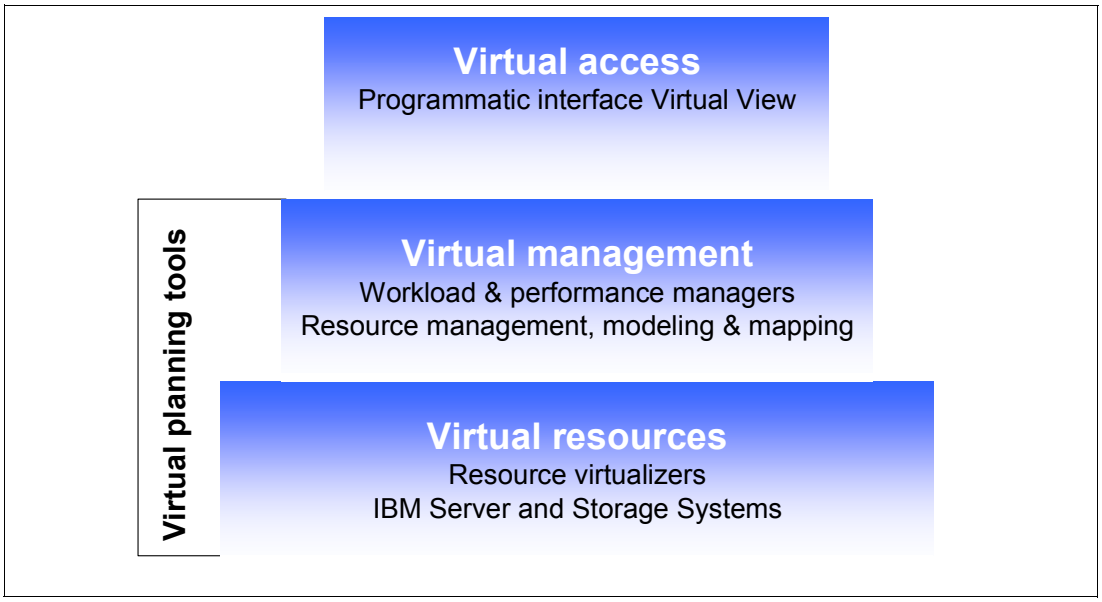

Figure 1-3 The IBM Virtualization Engine platform

Important: This book uses the following terminology:

- The *IBM Virtualization Engine platform* refers to the full virtualization capabilities that are offered by IBM, including hardware technologies and software components.
- The *IBM Virtualization Engine Systems Edition* refers to IBM @server technologies that are delivered through server and storage products. This term also includes resource virtualizers. It is described in Chapter 2, “The IBM Virtualization Engine Systems Edition” on page 21.
- The *IBM Virtualization Engine Management Collection* refers to the virtual management and virtual access capabilities that are delivered through software components, namely Virtualization Engine console, IBM Director, and Enterprise Workload Manager. It is described in Chapter 3, “The IBM Virtualization Engine Management Collection” on page 35. Resource Dependency Service (RDS) is a new offering of the Virtualization Engine Management collection. This redbook is about Version 1 of the Virtualization Engine, which does not include this offering.
- In the past, we referred to Virtualization Engine *Release*. However, today we use Virtualization Engine *Version*. Both refer to the same Virtualization Engine level.

What this Redbook is about:

- This redbook explains how to install the following Virtualization Engine virtual management and virtual access components: the Virtualization Engine console, the IBM Director Multiplatform, and the Enterprise Workload Manager.
- It includes the installation of the IBM systems provisioning component, the IBM Tivoli Provisioning Manager and how it can be used with the former Virtualization Engine platform components to help fix customer pain points.
- It does not detail how to set up the IBM Virtualization Engine Systems Edition.

Having the ability to efficiently manage the IT infrastructure is a journey, not a major re-engineering project, that requires massive rip-and-replace steps. There is a stepped approach that most successful companies follow, as shown in Figure 1-4.

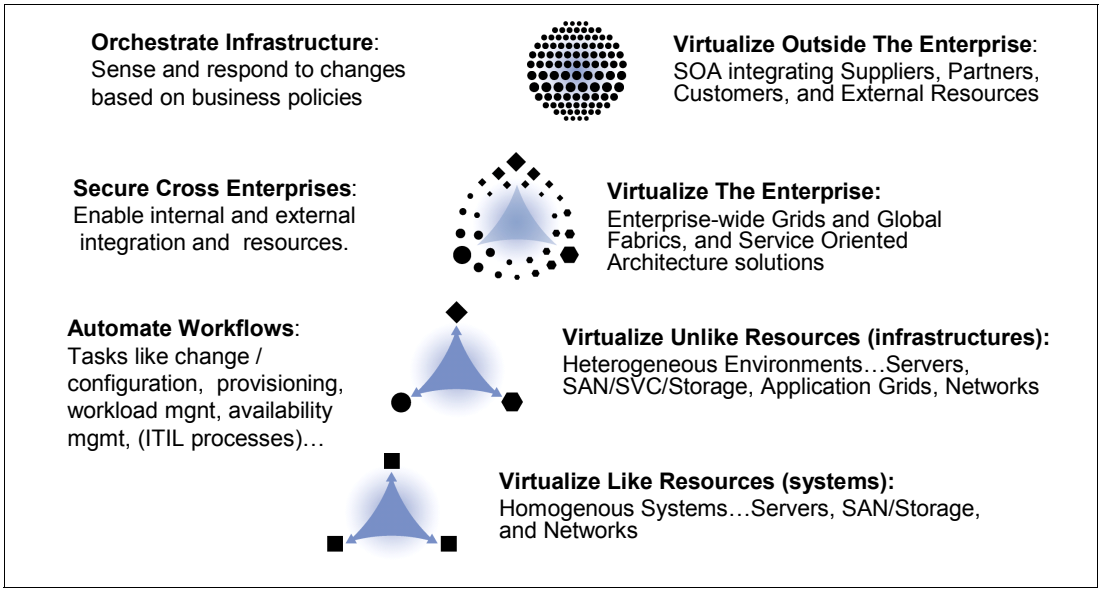
Orchestrate Infrastructure:
Sense and respond to changes based on business policies
Virtualize Outside The Enterprise:
SOA integrating Suppliers, Partners, Customers, and External Resources
Secure Cross Enterprises:
Enable internal and external integration and resources.
Virtualize The Enterprise:
Enterprise-wide Grids and Global Fabrics, and Service Oriented Architecture solutions
Automate Workflows:
Tasks like change / configuration, provisioning, workload mgnt, availability mgmt, (ITIL processes)…
Virtualize Unlike Resources (infrastructures):
Heterogeneous Environments…Servers, SAN/SVC/Storage, Application Grids, Networks
Virtualize Like Resources (systems):
Homogenous Systems…Servers, SAN/Storage, and Networks

Figure 1-4 Stages of virtualization solutions

First you simplify the environment by consolidating like systems onto fewer, more manageable resources. We supply the systems technology, together with products such as the IBM Director, to allow customers virtualize like resources. These are resources on homogenous systems, storage, and networks.

The next step is to automate those resources and add or move resources as required by the business. Allowing business needs to drive resources dictates how well the business performs. We also deliver on this second stage of virtualization through the System Services that allow you to virtualize unlike system resources, application based grids, and networks.

As the evolution of the IBM Virtualization Engine continues, more platforms are supported by the various virtualization solutions.

1.4 The value of the IBM Virtualization Engine

Altogether, provisioning and orchestration solutions, autonomic computing technologies, and virtualization facilities and services are able to simplify the IT infrastructure and build the core of an On Demand Operating Environment. More integration among the products that is based on these concepts is planned for the coming years.

Making a business more flexible results in a quicker response to changes such as customer demand, market conditions, competitive landscape, and mergers and acquisitions. The Virtualization Engine helps you:

- Manage the complexity of the infrastructure.
- Reduce the cost of your infrastructure management.
- Improve the utilization of your IT resources.
- Improve the quality of service that you offer to your customers.

1.4.1 Managing infrastructure complexity

When you simplify your IT structure, you can reduce costs and complexity. You can also improve reliability and prioritize IT and business decisions so that they better support your business. Virtualization can help you achieve these goals.

With solutions for virtualization you can:

- Lower the cost of your existing infrastructure.
- Reduce the complexity of managing and adding to your infrastructure.
- Enhance the flexibility of your infrastructure.

The IBM Virtualization Engine is a comprehensive portfolio of systems technologies and tools that can help you aggregate pools of resources and get a consolidated view of them throughout your IT environment. Virtualization Engine can help you automate the management of these resources based on your business goals and make basic systems management of multiple disparate systems possible. With IBM Virtualization Engine, you can start to put in place the technology that enables real-time, dynamic deployment and optimization of IT resources. Its open standards-based approach allows you to manage all your systems from a single, easy-to-use dashboard.

1.4.2 Simplifying management and reducing cost

The IBM Virtualization Engine provides a comprehensive set of services and technologies that can help you achieve your goals by aggregating pools of computing resources into a single, logical view. This can help to streamline IT administration, lower costs, and free skilled IT staff to work on other tasks.

A key enabler of simplification is consolidated systems management, which allows you to use a Web-based interface to perform core systems management functions in a multiple IBM server and OS environment. Managing your systems in a consistent way can help reduce your organization's total cost of ownership.

Using an open standards base as its framework, consolidated systems management — delivered through IBM Director Multiplatform and the Virtualization Engine console — makes it easier to implement a solution that works across multiple IBM platforms, helping you to integrate third-party products and to migrate smoothly to advanced autonomic and on demand features while you continue to leverage existing management investments.

1.4.3 Improving IT utilization

You can improve IT utilization with the simplified, cross-platform view that the IBM Virtualization Engine provides. With Virtualization Engine, you can:

- Increase use of system resources, thereby reducing IT proliferation.
- Broaden access to resources, which helps to give people access to the information and applications they need when they need them.
- Increase application availability to help operational speed and business responsiveness.
- Leverage flexible technologies for your changing business needs.

The IBM Virtualization Engine helps optimize both IBM systems and IBM TotalStorage systems across your organization. Because it can provide world-class virtualization functions across platforms — IBM systems and selected others — the IBM Virtualization Engine can enable simplified management of resources across a truly heterogeneous infrastructure.

In addition, you can leverage IBM Virtualization Engine's optimization, provisioning, and systems management functionality for one server, or you can extend it to all enabled IBM and TotalStorage systems across your entire enterprise. By monitoring workloads across your infrastructure, the EWLM component of Virtualization Engine gives you the ability to better use IT resources by helping increase utilization of existing systems and by employing EWLM-enabled routers that can direct workloads to resources programmed to meet service goals. EWLM can encompass a broad heterogeneous infrastructure, including network resources. Alternatively, you can deploy EWLM to only handle subsets of your IT infrastructure, depending on your current technology requirements.

1.4.4 Improving service levels

Continuity of business operations is a must for an on demand business. IT must not only fix problems when disruptions occur, but must also address potential threats proactively before they impact the business. Users and partners must be able to access the right systems securely, when they need them. Improved service levels let you:

- Adapt and reconfigure your infrastructure depending on business needs, which results in higher quality of service.
- Analyze operating environment events and, if delivery of critical services is threatened, you can rapidly detect and fix problems.
- Balance priorities by orchestrating execution of management processes.

2

The IBM Virtualization Engine Systems Edition

The Virtualization Engine delivers both systems services and systems technologies to help achieve virtualization in an IT environment. It provides the following advantages:

- Pools all computing resources (servers, storage, and network appliances, both IBM and others) into one virtual environment .
- Simplifies infrastructure and reduces management complexity by building integrated virtualization platform technologies into every box .
- Provides consistency across the IBM and TotalStorage brands.
- Uses a common open interface.
- Provides an integrated approach to enterprise-wide virtualization.

The IBM Virtualization Engine Systems Edition uses the IBM technologies that are integrated into the IBM @server family. These technologies are delivered within the IBM servers, including pSeries, iSeries, xSeries and BladeCenter®, and zSeries. Some of these technologies are identical or almost identical across the different platforms. For clarity, the descriptions of these technologies are not duplicated in this chapter. Only the most specific features are described.

© Copyright IBM Corp. 2005. All rights reserved.

2.1 Virtualization technologies for p5 and pSeries

The virtualization technologies for the p5 and pSeries include the following:

- Dynamic LPAR
- Advanced Power Virtualization (Virtual I/O, Partition Load Manager, Micro-Partitioning™)
- IBM Director
- Cluster Management

Some of these functions are delivered through the Advanced POWER™ Virtualization feature (APV). APV is a combination of hardware enablement that includes components that are available together as a single priced feature. These feature allow three Virtualization Engine systems technologies to be ordered:

- Firmware key, which enables Micro-Partitioning, the shared processor pool, and Reserve CoD (p5-550 and p5-570 only). For initial system orders, the firmware is shipped activated. For upgrade orders, IBM ships a key that enables activation through an HMC panel.
- Virtual I/O Server, software which runs in a *host partition* to enable virtual SCSI for sharing Fibre Channel and SCSI adapters and the attached disks and virtual networking to enable sharing of Ethernet adapters. The VIO installation media ships with the above feature codes on POWER5 systems and supports NIMOL (from HMC) or CD installation. It supports AIX 5L™ V5.3 and Linux Virtual I/O clients.
- Partition Load Manager (PLM), which is a software tool for managing processor and memory resources automatically among partitions that are running AIX 5L V5.3 and AIX 5L V5.2 with the 5200-04 Recommended Maintenance Package.

You can find more information about the APV feature in *Advanced POWER Virtualization on IBM p5 Servers:Introduction and Basic Configuration*, SG24-7940.

2.1.1 Dynamic logical partitioning (DLPAR) and Micro-Partitioning

Logical partitions allow you to distribute resources within a single IBM @server p5 server to allow it to function as though it were multiple independent servers. Logical partitions use dedicated and shared processors to run concurrent operations. System resources (processors, memory, and I/O components) can be added and deleted from partitions while they are active. Dynamic logical partitioning allows you to move resources to, from, and between running logical partitions without shutting down or restarting the logical partitions. The ability to reconfigure DLPARs enables system administrators to dynamically redefine available system resources to reach optimum capacity for each partition.

AIX, i5/OS™, and Linux are supported operating systems that can be installed on IBM @server p5 servers. These operating systems operate as independent logical servers. However logical partitions share a few system attributes, such as the system serial number, system model, and processor feature code. All other system attributes might vary among logical partitions. The server hypervisor directly controls the resource allocations on the server and the communications between logical partitions on the server.

The Hardware Management Console (HMC) is a hardware appliance that connects to the server. The HMC is used to specify to the server firmware how you want resources to be allocated among the logical partitions on the managed system. The HMC is also used to start and stop the logical partitions, to update the server firmware code, and to transmit service information to service and support if there are any hardware problems with your managed system. One HMC is required to configure and manage logical partitions

The POWER5 processor introduces an enhanced partitioning model based on established mainframe technologies and LPAR or DLPAR implementations on POWER4™ and POWER4+™ servers. Micro-Partitioning enables the virtualization of system resources on an extremely granular level. It gives you the ability to assign fractions of a processor to a logical partition. Before micro-partitioning, the system could use only whole processors. However, the POWER Hypervisor allows you to specify the processing power for a logical partition in hundredths of a processor, and you can create logical partitions in increments of one tenth of a processor. The value of micro-partitioning is that partitions, whether specified as less than a processor or more, can be created with more precisely match the workload requirements.

You can still assign whole numbers of dedicated processors to a logical partition. All processors that are not dedicated to individual logical partitions are put into a shared processor pool. The POWER Hypervisor allows logical partitions to use the unused processing power of other logical partitions when needed.

Multiple tools are available to help you plan and deploy the most optimized configuration. For example, the LPAR Validation Tool (LVT) assists the user in the design of LPAR systems and provides an LPAR validation report that reflects the user's system requirements while not exceeding LPAR recommendations. The LVT is a PC based tool you can dowload from the following Web site:

http://www-03.ibm.com/servers/eserver/iseries/lpar/systemdesign.html

To know more about partitioning in an AIX environment, refer to the following Web sites:

- http://publib.boulder.ibm.com/infocenter/eserver/v1r2s/en_US/info/iphbk/iphbk.pdf
- http://www-1.ibm.com/servers/eserver/pseries/lpar/
- http://publib.boulder.ibm.com/infocenter/pseries/index.jsp?topic=/com.ibm.help.doc/welcome.htm

You can also refer to *Partitioning Implementation for IBM p5 Servers*, SG24-7039.

2.1.2 IBM Cluster Systems Management

While Cluster Systems Management (CSM) provides basic cluster management functions such as hardware control and operating system and software installation, it also provides significant value-added functions such as configuration file management, a robust monitoring infrastructure, automated event management, diagnostic probes, and common management of AIX and Linux clusters.

2.1.3 Partition Load Manager

PLM is a resource manager that assigns and moves resources based on defined policies and utilization of these resources. PLM manages memory, dedicated and shared processors partitions and use Micro-Partitioning technology to readjust resources. This adds additional flexibility on top of the micro-partitions flexibility added by the POWER Hypervisor. PLM requires POWER5 hardware and requires DLPAR capable partitions.

PLM is installed as a licensed program product in a partition running AIX (AIX 5L V5.2 ML4 or AIX 5L V5.3 are required). PLM does not support Linux or i5/OS clients. One instance of PLM can manage one single server and supports a group of partitions, called managed partition group defined in a policy. The partition running PLM does not have to be dedicated.

I/O resources are not managed by PLM. PLM has no knowledge about the importance of a workload running in the partitions, thus it is not able to adjust priority based on the type of the workload.

The PLM application is based on a client/server model for the sharing of system information, such as processor or memory events, across concurrent present LPARs. The managed partitions notify the PLM server when resources are under or over utilized (according to predefined thresholds) and then the PLM server makes resources allocation decisions based on a policy defined by an administrator. PLM needs to communicate with the HMC to gather system information.

Every node managed by PLM must be defined in the policy in which the thresholds and the partition priority are defined. A partition cannot fall below its minimum or rise above its maximum definition for each resource. A partition priority cannot be changed upon the arrival of another priority work. The priority can be changed by loading a new policy.

When a request for a resource is received, PLM takes resources in the following priority order: from unallocated resources, from partitions which have not reach a lower defined threshold, or from a lower priority partition with excess resources. If the request for a resource cannot be honored, it is queued an reevaluated later.

You can learn more about PLM by reading *Advanced POWER Virtualization on IBM p5 Servers:Introduction and Basic Configuration*, SG24-7940.

2.1.4 Virtual I/O Server

The Virtual I/O (VIO) Server is a special-purpose partition that provides I/O resources capability to client logical partitions. VIO Server is a single function appliance that resides in a pSeries POWER5 partition that:

- Allows creation of partitions without requiring additional physical I/O resources.
- Enables creation of more partitions than I/O slots or physical devices. Partitions can have dedicated I/O, Virtual I/O (VIO), or both.
- Provides virtual SCSI and Shared Ethernet Adapter function to client partitions.

Virtual I/O is included in the Advanced POWER Virtualization feature for IBM @server p5 systems and requires AIX 5L V5.3 or SuSE LINUX Enterprise Server 9 for POWER (SLES 9) or Red Hat Enterprise Linux AS 3 for POWER, Update 3 (RHEL AS 3).

There are two different versions of VIO Server:

- APV Version 5765-G34 - For POWER5 Systems.
- AOPV Version 5724-K81 - For OpenPower™ Systems

The VIO Server is a feature that enables partitions to share I/O resources such as Ethernet adapters and SCSI or Fibre Channel disks. The VIO Server is a special POWER5 partition that owns and provides access to the real physical resources that are shared with clients. Operating system running in a logical partition thinks it has direct access to a disk or a network, though that access is in fact executed through the VIO Server. From the application perspective, VIO looks similar to physical I/O.

A physical adapter assigned to a VIO Server partition can be shared by one or more partitions, enabling administrators to minimize the number of physical adapters they require for individual clients. The VIO Server can thus reduce costs by eliminating the need for dedicated network adapters, disk adapters and disk drives.

In order to use VIO, one VIO Server needs to be created using the HMC. This is a special type of partition that runs a specific kernel intended for the purpose of providing VIO. Applications cannot run in this partition. Multiple VIO Servers can be defined for availability or performance reason.

There are two types of Virtual I/O provided on IBM @server p5 systems: the virtual SCSI and the virtual Ethernet adapters.

Virtual SCSI

Virtual SCSI adapters provide access to virtual disks, physical logical unit numbers (LUN) and disk drives. A virtual disk does not have to map one-to-one with a physical disk. The physical disk drive is owned by a Virtual /IO Server and can be a parallel SCSI or Fibre Channel-attached disk. It can be a LUN on an advanced function controller such as the IBM TotalStorage Enterprise Storage Server®. It can also be a virtualized disk on a SAN Virtualization Controller.

To the operating system, a virtual disk looks similar to a physical disk. Usual tools and products, such as file systems or data backup and restore utilities can be used against these disks. The value of a virtual disk is that it can be sized to match the requirements of the operating system.

Virtual Ethernet

Virtual Ethernet is an internal LAN that connects to a set of partitions. These LANs exist only in the memory of the server and are implemented via the POWER Hypervisor. They provide fast communication between partitions and look similar to Ethernet LANs to the operating system in the partition. If these LANs need to connect to external physical LANs, an Ethernet adapter is needed to connect to the virtual LAN. A Shared Ethernet Adapter allows multiple virtual LANs to connect to the external network via the same physical Ethernet Adapter.

Without requiring any additional hardware, the POWER5-based p5 systems provide high-speed Virtual Ethernet communication paths between multiple operating systems such as AIX 5L and Linux. Virtual Ethernet segments can be dynamically created and access to a virtual LAN segment can be restricted for security or traffic segregation requirements.

2.2 Virtualization technologies for iSeries

The virtualization technologies for the iSeries include the following:

- Dynamic LPAR (DLPAR)
- Virtual Ethernet
- Virtual I/O
- Micro-partitioning
- IBM Director
- Capacity on Demand (CUoD)
- Virtualization Engine console

Some of the them are described in the following sections. You can find more information in *IBM @server* i5 and iSeries System Handbook:IBM i5/OS Version 5 Release 3, GA19-5486-26.

Furthermore, IBM Virtualization Engine for iSeries (5733-VE1) includes the IBM Director Multiplatform on all V5R3 iSeries operating system orders and IBM Enterprise Workload Manager on selected iSeries editions. The Virtualization Engine console is included.

2.2.1 Dynamic logical partitioning

Dynamic logical partitioning has been introduced in 2.1.1, “Dynamic logical partitioning (DLPAR) and Micro-Partitioning” on page 22. Logical partitioning is required for i5/OS on

IBM @server p5 servers. The maximum number of partitions supported depends on the server model and the HMC version used to partition the server. Using i5/OS V5R3:

- IBM @server i5 models support up to 64 i5/OS logical partitions and up to 254 total logical partitions, with each partition requiring a minimum of 0.10 processing units.
- IBM @server p5 models that support i5/OS support up to 20 i5/OS logical partitions, with each logical partition requiring a minimum of 0.10 processing units.

The total processing capability that is dedicated to i5/OS logical partitions on IBM @server p5 models varies by server model but is not to exceed two processors.

The HMC is required to partition the server.

The processors can be changed dynamically without restarting the partition. They can be shared among multiple partitions in shared mode of capped and uncapped. Powered-off partitions using dedicated processors have processors available to shared processor pool.

The memory can be changed dynamically without a restart of the partition and can be assigned in increments of 16 MB, 32 MB, 128 MB, 164 MB, and 256 MB.

IOP can be switched among partitions dynamically. I/O resources are allocated at the slot level.

The Logical Partition Validation Tool (LVT) is a tool to assist you for planning and validating the logical partition configuration. The LVT is a PC based tool intended to be run as a standalone Java application. You can download it from:

```
//ftp.software.ibm.com/as400/web/lpar/LVT4_0504_JRE_Build.exe
```

For more information, see:

http://publib.boulder.ibm.com/infocenter/eserver/v1r2s/en_US/info/iphbl/iphbl.pdf

2.2.2 Micro-Partitioning

The POWER5 processor introduces an enhanced partitioning model based on established mainframe technologies and LPAR or DLPAR implementations on POWER4 and POWER4+ servers. Micro-Partitioning enables the virtualization of system resources on an extremely granular level. In POWER5 processor-based systems, physical resources can be abstracted into virtual resources that are available to partitions. Resources can be shared easily, and changes in resource allocation are transparent to users.

2.2.3 Virtual Partition Manager

The Virtual Partition Manager (VPM) introduces the capability on the iSeries to create and manage Linux partitions without the use of the HMC. The Virtual Partition Manager supports the needs for customers that want to add simple Linux workloads to their IBM @server i5 server. The VPM is included with i5/OS 5R3. The VPM supports environments with one i5/OS partition and up to four Linux partitions. In addition, the single i5/OS partition must own and manage all of the I/O resources.

For more information, refer to *Virtual Partition Manager: a Guide to Planning and Implementation*, REDP-4013.

2.2.4 Virtual I/O

Virtual I/O (VIO) adapters can be defined in a partition profile so that they are added to the logical partition when the logical partition is activated. VIO adapters can also be added and removed dynamically. You can change the properties of a VIO adapter without requiring a restart of the logical partition. i5/OS logical partitions support up to 32 767 VIO adapters.

However, VIO adapters cannot be moved dynamically between logical partitions: after the logical partition starts, you cannot use dynamic logical partitioning to remove the required I/O device from the running logical partition or move the required I/O device to another logical partition.

Virtual SCSI adapters, virtual serial adapters, and virtual Ethernet adapters are supported for i5/OS logical partitions.

Virtual Ethernet

Without requiring any additional hardware, the POWER5-based iSeries systems provide 1 Gb Virtual Ethernet communication paths between multiple operating systems such as i5/OS, Linux and AIX 5L. Virtual Ethernet segments can be created dynamically and access to a virtual LAN segment can be restricted for security or traffic segregation requirements. Up to 4094 networks can be defined.

Virtual OptiConnect

The Virtual OptiConnect feature emulates external OptiConnect hardware by providing a virtual bus between logical partitions. You can use the Virtual OptiConnect feature without any additional hardware requirements. To use Virtual OptiConnect, you only need to install OptiConnect for OS/400® (a priced optional feature).

2.3 Virtualization technologies for zSeries

The virtualization technologies for the zSeries include the following:

- Dynamic LPAR (DLPAR)
- Parallel Sysplex and GDPS® Clustering
- Virtual I/O
- Intelligent Resource Director (IRD) / WLM
- z/VM
- Virtual Machine Resource Manager
- Guest LANs
- IBM Director
- Hipersockets
- Capacity on Demand (CUoD)
- zAAP

2.3.1 Virtual servers: LPAR and z/VM

Dynamic logical partitioning increases flexibility, which enables selected system resources (such as processors, memory, and I/O components) to be added to and deleted from dedicated partitions while they are actively in use. The ability to reconfigure dynamic LPARs enables system administrators to redefine all available system resources dynamically to reach optimum capacity for each partition.

Today, zSeries boasts up to 30 logical partitions with PR/SM™, and can deploy thousands of virtual servers when using z/VM. All of which can share the physical resources: processors,

memory, and I/O. zSeries virtualization technology can support Linux, z/OS, z/VM, Transaction Processing Facility (TPF), and z/VE with a mixed operating systems environment.

2.3.2 Parallel Sysplex clustering

A Parallel Sysplex cluster consists of up to 32 z/OS images, coupled to one or more Coupling Facilities (CF) using high-speed specialized links for communication. The Coupling Facilities, at the heart of the Parallel Sysplex cluster, enable high-speed, record-level read/write data sharing among the images in a cluster. A Parallel Sysplex cluster has no single point of failure and can provide customers with near-continuous application availability over planned and unplanned outages.

Running a zSeries Parallel Sysplex can help to simplify a distributed environment by running the transaction volumes on a single unified processing base with the associated benefits of zSeries availability. The Parallel Sysplex ability to support very large databases with integrity and high performance can help simplify those environments which have to partition data across multiple servers. This can help increase availability while reducing administrative and management costs.

2.3.3 Workload Manager for z/OS

The workload manager (WLM) for z/OS provides the functions to define, to implement, and to monitor system performance against business goals.

The workload manager for z/OS, introduced in 1994 into the MVS™ (and now called the z/OS) Operating System, wanted to provide help for balancing and distributing workloads which compete for processor resources, such as CPU and storage, inside a server by providing performance and topology information. Work is defined in terms of business service level goals, taking into account the response time needs of customer business applications, together with the relative importance of all the work to be managed. Workload management introduced a sysplex-wide view of performance administration, monitoring, and management of workload running under z/OS.

The WLM algorithms work in cooperation with subsystem work managers to check if the performance goals are being met and to adapt the resources accordingly. WLM introduced the concept of the *service definition*, that contains all the information needed for WLM processing (service class, service policy, importance of the goal for the business, classification rules, and so forth).

WLM facilities continue to be refined, upgraded and enhanced in an effort to improve resource and performance management and reporting capabilities, as well as providing support for new technologies.

You can find more information in the following:

- *z/OS Version 1 Release 6 Implementation*, SG24-6377
- *OS/390 Workload Manager Implementation and Exploitation*, SG24-5326
- *System Programmer's Guide to: Workload Manager*, SG24-6472

2.3.4 HiperSockets

zSeries HiperSockets is a technology that provides high-speed TCP/IP connectivity between virtual servers running within different logical partitions (LPARs) of a zSeries server. It eliminates the need for any physical cabling or external networking connections between these virtual servers.

HiperSockets is a totally integrated, any-to-any, virtual TCP/IP network that provides benefits not achievable by grouping servers around a mainframe interconnected by external networking technology. It can provide the ever-available, high speed network connection among multiple combinations of virtual servers and LPARs running Linux on zSeries, Linux on S/390®, z/OS V1R2, z/VM V4R2 and VSE/ESA™ V2R7 or later on zSeries.

The HiperSockets function provides users with attachment to high-speed *logical* LANs with minimal system and network overhead. HiperSockets provides the fastest TCP/IP communication between consolidated Linux, z/VM, VSE/ESA and z/OS virtual servers on a zSeries server. It provides internal *virtual* Local Area Networks which act similar to TCP/IP networks within the zSeries server. This integrated zSeries Licensed Internal Code (LIC) function, coupled with supporting operating system device drivers, establishes a higher level of network availability, security, simplicity, performance, and cost effectiveness than is available when connecting single servers or LPARs together using an external TCP/IP network.

HiperSockets eliminates the need to use I/O subsystem operations when communicating between LPARs in the same zSeries server. It offers significant value in server consolidation connecting many virtual servers. It does not use an external network and eliminates attachment costs while improving availability and performance.

For more information, refer to *IBM @server* zSeries Connectivity Handbook, SG24-5444.

2.3.5 Intelligent Resource Director

The Intelligent Resource Director (IRD), introduced as part of the IBM zSeries and z/OS, is designed to give your installation an enhanced ability to dynamically move resources to your most important work. IRD can help you handle unexpected workload spikes and improve your system's efficiency and availability. It also can help reduce the systems management skills and time required to define the I/O configuration.

Intelligent Resource Director uses facilities in z/OS Workload Manager (WLM), Parallel Sysplex, and PR/SM. It helps you make sure that all those resources are being utilized by the right workloads, even if the workloads exist in different Logical Partitions (LPs).

IRD is a feature that extends the concept of goal-oriented resource management by allowing you to group system images that are resident on the same zSeries server running in LPAR mode, and in the same Parallel Sysplex, into an *LPAR cluster*. This gives WLM the ability to manage resources, both processor and I/O, not just in one single image but across the entire cluster of system images.

Through IRD, WLM adjusts the number of logical processors within an LPAR and the processor weight dynamically based on the WLM policy. The ability to move the CPU weights across an LPAR cluster provides processing power to where it is most needed based on WLM goal mode policy. Through IRD, the dynamic channel path management is able to change channel bandwidth between disk control units to address current processing needs. IRD allows the priority queuing of I/O requests in the channel subsystem and the specification of relative priority among LPARs.

For more information about IRD, refer to the following:

- *z/OS Intelligent Resource Director*, SG24-5952
- *IBM @server* zSeries 900 Technical Guide, SG24-5975

2.3.6 Virtual local area network

A virtual local area network (VLAN) is a group of workstations with a common set of requirements that are independent of physical location. VLANs have the same attributes as a physical LAN, even though they cannot be located physically on the same LAN segment. VLANs can be used to increase bandwidth and reduce overhead by allowing networks to be organized for optimum traffic flow.

VLAN facilitates easy administration of logical groups of stations that can communicate as though they were on the same LAN. They also facilitate easier administration of moves, adds, and changes in members of these groups. VLANs are also designed to provide a degree of low-level security by restricting direct contact with a server to only the set of stations that comprise the VLAN.

z/OS V1R4 or later releases provide full VLAN. This support is also applicable to the Linux environment. This support can depend on the level of software and feature codes.

2.3.7 zSeries Application Assist Processor

With the zSeries Application Assist Processor (zAAP), which is available on the latest IBM @server zSeries, users can deploy and integrate new Java technology-based workloads on the very same platform as heritage applications and core business databases in a highly cost-effective manner.

The zAAP optional assist feature allows customers to purchase additional processing power exclusively for Java application execution without affecting the total Million Service Units (MSU) rating or machine model designation. zAAPs are designed to operate asynchronously with the general CPs to execute Java programming under control of the IBM Java Virtual Machine (JVM™). zAAPs can only help execute Java applications and application servers that use the IBM JVM. The IBM JVM processing cycles can be executed on the configured zAAPs with no modifications to the Java applications.

zAAPs enable you to:

- Simplify and reduce server infrastructures by integrating e-business Java Web applications next to mission critical data for high performance, reliability, availability, and security.
- Maximize the value of your zSeries investment through increased system productivity.
- Lower the overall cost of computing for WebSphere Application Server and other Java technology-based applications through hardware, software, and maintenance savings.

For more information, refer to the following:

- *Effective zSeries Performance Monitoring Using Resource Measurement Facility*, SG24-6645
- *zSeries Application Assist Processor (zAAP) Implementation*, SG24-6386

2.4 Virtualization technologies for xSeries and BladeCenter

The virtualization technologies for the xSeries and BladeCenter include the following:

- Integrated shared infrastructure for Blades
- IBM Director
- VMWare
- Virtual Machine Manager
- Capacity Manager
- Integrated shared infrastructure

2.4.1 IBM Director

IBM Director is a comprehensive systems-management solution and a powerful suite of tools and utilities. IBM Director automates many of the processes required to manage systems proactively, including preventive maintenance, diagnostic monitoring, troubleshooting, and more. It offers a graphical user interface that provides system administrators easy access to both local and remote systems. IBM Director can be used in environments with multiple operating systems (heterogeneous environments). IBM Director enables monitoring and event management across a heterogeneous IT environment, including Intel, Windows, Linux, AIX and i5/OS, from a single Java-based user interface. From one access point, you can monitor system resources, inventory, events, task management, core corrective actions, distributed commands and hardware control for your servers and storage.

With IBM Director, IT administrators can view and track the hardware configuration of remote systems in detail and monitor the usage and performance of critical components, such as processors, disks, and memory.

With IBM Director, you have complete access to the systems management hardware in addition to other management tasks such as event management, inventory, and deployment. The BladeCenter Assistant of IBM Director has some additional BladeCenter-specific tasks and can act as an alert destination for xSeries service processors and the management module. The IBM Director Console allows system administrators to manage multiple BladeCenter chassis from a common interface.

From the IBM Director Console, system administrators can monitor resource utilization for performance and capacity planning or can check errors such as a hardware failure. IBM Director also allows you to reconcile and install software remotely from the console interface, which can be useful in large environments.

Attention: All the IBM Director functions are not available on all the platforms. They are described in detail in each platform documentation.

2.4.2 VMware

VMware is a partitioning software designed for server consolidation, rapid deployment of new servers, availability, and simplified management. It helps to improve hardware utilization, to save space, and to reduce or to free up IT staffing and hardware costs. Using VMware, you can build a pool of logical computing resources from the physical servers. You can partition physical servers into secure virtual machine servers. You can then isolate operating systems and applications in the multiple virtual machine servers that reside on a single piece of hardware. Then, you can distribute the resources to any operating system or application as needed, when needed. The VMware virtualization layer maps the physical hardware resources to the virtual machine's resources, so each virtual machine has its own CPU, memory, disks, and I/O devices.

There are three versions of VMware:

- *VMware Workstation* allows users to run multiple x86-based operating systems, including Windows, Linux, and NetWare, and their applications simultaneously on a single PC in fully networked, portable virtual machines
- *VMware GSX Server* allows virtual machines to be managed remotely, provisioned automatically, and standardized on a secure, uniform platform. VMware GSX Server transforms physical computers into a pool of virtual machines. GSX version is an application that is loaded on top of Windows 2000 Server. GSX is a Windows application supported by VMWare, not IBM.
- *VMware ESX Server* is a virtual infrastructure software for partitioning, consolidating, and managing systems in mission-critical environments. ESX Server and VMware Virtual Infrastructure Nodes provide a highly scalable virtual machine platform with advanced resource management capabilities, which can be managed by VMware VirtualCenter. For example, VMware ESX Server could migrate, on-the-fly, an actual running VMware image from one ESX server to another ESX server.

  ESX version is based on Linux. VMware ESX Server allows you to:

  - Implement server consolidation: applications and infrastructure services can be consolidated onto fewer highly scalable, highly reliable enterprise-class servers.
  - Deliver high availability: with a stable, uniform platform, you can deploy new solutions faster and more efficiently.
  - Guarantee service levels: such as an internal service provider, your IT department can deliver guaranteed server resources of CPU, memory, disk bandwidth, and network bandwidth at optimum performance levels, improving service to customers.
  - Streamline development and testing: your software developers can work with multiple machine environments and build more realistic tests in less time with less hardware.
  - Remotely manage servers from any location, simplifying server maintenance

  VMware ESX Server runs directly on your system hardware to provide a secure, uniform platform for deploying, managing, and remotely controlling multiple operating systems. VMware ESX Server V2.x consists of a base product and optional add-on components.

Different VMware components offer multiple functions, in particular the VirtualCenter, the P2V assistant and Vmotion:

- The VMware VirtualCenter is a virtual infrastructure management software that centrally manages an enterprise's VMs as a single, logical pool of resources. VirtualCenter provides the central point of control for workload management, provisioning and availability. The VirtualCenter interface provides an overview of all the virtual resources in a data center. The VirtualCenter server can execute user-specified scheduled tasks automatically, such as powering on or moving powered-off virtual machines.
- The VMware P2V Assistant Version 2.0 is an enterprise-class migration tool that transforms an image of an existing physical system into a VMware virtual machine. This tool enables fast and reliable physical to virtual machine migration.
- VMotion allows you to move a powered on virtual machine between hosts in a farm. Moving a powered on virtual machine allows the virtual machine to continue performing transactions without interruption. The benefits of VMotion in a blade server environment are:

  - Improves overall utilization across the data center.
  - Reacts more quickly to changing business demands.
  - Creates alternate form-factors for cost-effective consolidation.

VMware ESX Server and VMware VirtualCenter are available through IBM xSeries channels as standard xSeries part numbers.[1]

For more information, read the following:

- *Tuning IBM @server* xSeries Servers for Performance, SG24-5287
- *VMware ESX Server: Scale Up or Scale Out?*, REDP-3953
- *Implementing VMware ESX Server 2.1 with IBM TotalStorage FAStT*, SG24-6434

2.4.3 Virtual Machine Manager

IBM Virtual Machine Manager (VMM) is an extension to IBM Director that allows you to manage both physical and virtual machines from a single console. VMM is vendor agnostic; with VMM, you can manage both VMware ESX Server and Microsoft Virtual Server environments using IBM Director. VMM also integrates VMware VirtualCenter and IBM Director for advanced virtual machine management.

VMM allows you to:

- Provide single glass management of both physical and virtual systems.
- See relationships between physical and VMM logical objects.
- Review status of physical hosts and virtual machines.
- See events that have occurred on the physical host or virtual machine.
- Move running virtual machines manually or automatically between two physical hosts.
- Perform predictive failure analysis, virtual machine migration, and provisioning.

[1] Under a joint global distribution and support agreement, IBM can resell VMware ESX Server, VM software on IBM @server xSeries and BladeCenter systems.

3

The IBM Virtualization Engine Management Collection

The IBM Virtualization Engine management collection provides multiplatform system tools that help:

- Simplify management and use of IT resources.
- IT infrastructures support business process change.

The IBM Virtualization Engine Version 1 management collection includes the following components to enable enterprise-wide virtualization:

- IBM Virtualization Engine console
- IBM Enterprise Workload Manager
- IBM Director Multiplatform

The following components are part of the existing IBM Virtualization Engine Version 1 current offering, although they are not integrated into the IBM Virtualization Engine management collection:

- IBM Tivoli provisioning Manager
- IBM Grid Toolbox V3 for Multiplatforms

The IBM Virtualization Engine management collection provides a new offering in Version 2, the Resource Dependency Service.[1]

This chapter describes all of these components.

[1] IBM plans are subject to change without notice and this information represent goals and objectives only. Such information is not intended as a definitive statement of a commitment to specific levels of function or delivery schedules.

© Copyright IBM Corp. 2005. All rights reserved.

3.1 IBM Enterprise Workload Manager

The IBM Enterprise Workload Manager (EWLM) enables you to continuously monitor multi-tiered, distributed, heterogeneous workloads across an IT infrastructure consisting of multiple types of operating systems, to better achieve defined business goals for user services. These capabilities today allow you to:

- Identify work requests based on service class definitions.
- Track performance of those requests across server and subsystem boundaries.
- Redirect work, using EWLM-enabled routers, to achieve specific performance goals for each service class.

In the future, we expect that EWLM will have the capability to monitor and manage physical and virtual resources to achieve business goals across the full enterprise application portfolio. EWLM is intended to deliver levels of performance management not previously achievable for distributed systems. Its management reach can encompass your entire infrastructure, extending even to Internet servers.

EWLM R1.0, available today, allows you to set response time and business goals for user service and provides recommendation to take actions such as routing so that they are met. It can make possible cross-platform performance analysis and capacity planning. EWLM R1.0 monitors performance of an end-to-end application that spans hundreds of servers and operating systems. It can also manage network traffic routing for improved application performance and server effectiveness, according to the business needs expressed through a service class objective and an importance property.

3.1.1 The EWLM topology

A general overview of EWLM, as depicted in Figure 3-1 on page 37, introduces the following major components:

- The group of operating systems within the scope of an EWLM domain policy is collectively referred to as an *EWLM Management Domain*.
- An EWLM management domain has an associated EWLM Domain Manager, which is the focal control point for the entire domain and builds the global view. The EWLM Domain Manager is a Java-based server application which maintains the complete administrative view of policy, and coordinates administrative and reporting actions across all servers participating in the domain.

  The Domain Manager provides two processes:

  - One supports the Control Center.
  - One coordinates policy across the managed servers and aggregates performance statistics.

  The managed servers and the domain manager communicate via a lightweight publish and subscribe messaging protocol. This communication can be secured using SSL.

  The Domain Manager ensures global coordination of policy actions initiated from the Control Center. When a policy is activated, the Domain Manager sends it to all the managed servers in the management domain and begins collecting data from the managed servers in the context of the new policy. The Domain Manager aggregates the data that it collects and then stores it in an internal database for historical purposes. This data can be viewed in the Control Center to monitor how well the business goals are being met and to understand the topology of transactions, applications, and servers in the management domain. The Domain Manager keeps data collected from the managed

servers. In EWLM R1.0 the data can be displayed for as little as one minute or up to an hour interval.

- Administrative and operational interactions occur through a Web-enabled user interface, the Control Center. The Control Center provides the ability to administer, operate, and analyze the performance characteristics of the set of server operating system instances within an EWLM management domain (the managed servers). The Control Center is the focal point for:
  - The creation and the manipulation of the policy information.
  - The operational interactions to activate the policies, and to display the status of these policies and participating servers.
  - The reporting of performance statistics, and the drill-down into more detailed analytical data.
- Each managed server in the domain contains an EWLM Managed Server layer. This layer interacts with the underlying operating system, sending and receiving information to and from the EWLM Domain Manager. The collecting of transaction performance statistics and application topology data is based upon instrumentation of middleware application code, using the ARM 4.0 standard interface.

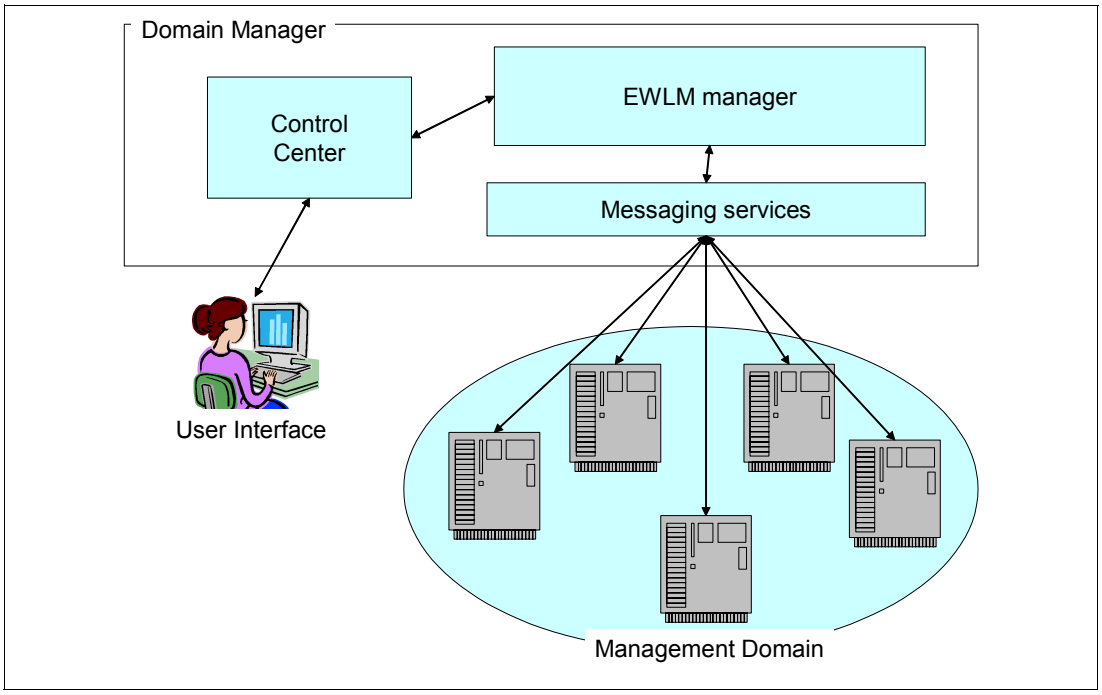

Figure 3-1 EWLM topology overview

3.1.2 The domain policy

The EWLM components are dependent upon policy definition, created through the user interface, that identifies the set of service classes, the performance objective for each service class, and the means by which each transaction instance is associated with a specific service class. The complete set of definitions is named a domain policy, which is represented in an XML document and which is installed into the EWLM Domain Manager.

An EWLM domain policy describes how the work can be classified, prioritized and monitored. A domain policy document contains the set of service class, transaction class, and optionally platforms and process classes. Only one service policy can be active at a given time in the

domain. The domain policy contains *performance goals* (for example average response time or percentile response time) and importance for each work.

The following performance goals are supported:

- Percentile response time: specifies what percentile of work requests should complete in a specified amount of time, specified as a percentage of the time.
- Average response time: specifies the average amount of time in which work in the service class should complete.
- Velocity: defines how fast work should run when ready, without being delayed for processor, storage and I/O (for managed system resources). Velocity goals are intended for work for which response time goals are meaningless, particularly for all kinds of processes which are not instrumented such as server processes, daemons or long running batch-like work. There are five categories of velocity, from the fastest to the slowest.
- Discretionary: for low priority work for which you do not have any particular performance goal. EWLM processes discretionary work using resources not required to meet the goals of other service classes.

For each service class, you then assign an *importance* indicating how important it is to your business that the service class performance goal be achieved. In the future we expect that EWLM will use the business importance to determine which work should give up resources and which work should receive more. Importance plays a role only when a service class is not meeting its goal. There are five levels of importance, from the highest to the lowest importance.

EWLM computes a value called a Performance Index (PI), to indicate how well a service class is attaining the goal defined for it. The Performance Index is a ratio between the response time goal and the actual response time. A PI greater than 1 means that the goal is being missed, while a PI less than 1 means that the goal is being met. The PI can be displayed for information in different output graphs.

3.1.3 The Application Response Measurement standard

In order for EWLM to report on transactions and potentially manage the domain based on these statistics, Application Response Measurement 4.0 (ARM) must be implemented for application instrumentation. ARM is an Open Group standard composed of a set of APIs that can be used to collect response time information for enterprise applications.

Applications that use ARM define transactions that are meaningful within the business process. Applications call ARM when transactions start or stop. The managed server in turn communicates with management applications which provide analysis and reporting of the data. The management server collects the status and response time. An important capability of ARM API and the ARM agent is the tracking of hierarchical relationships among transactions that span across multiple servers or operating systems. The ability to associate end-to-end transactions with units of work depends on the passing of a trace object on ARM calls. Knowledge of the parent-child relationships among transactions and the response times for each transaction enables an administrator to determine which transactions are delaying other transactions.

EWLM, ARM, and instrumented applications work together to provide end-to-end response times and application or transaction topology using correlation. Each EWLM Managed Server accumulates transaction class response time and topology information driven by ARM-instrumented applications. The EWLM Managed Server component extracts relevant resource consumption and delay statistics from the underlying operating system.

There is no need to change the application to use ARM. Middlewares and the operating systems, however, need to be ARM instrumented (having the ability to understand these ARM APIs) to report to EWLM. For the complete list of ARM-enabled products, see the following:

http://www-1.ibm.com/servers/eserver/about/virtualization/technical/faqs.html#ewlm

At the time of writing, the following products understand the ARM APIs:

- All the IBM operating systems are ARM-enabled. A layer of platform-dependent logic is provided, which is called Operating Systems extensions, that is tailored to the needs and capabilities of each operating system. The EWLM Managed Server component, a Java application, provides native interface classes that interact with the services provided by each operating system platform.
- DB2® Universal Database™, Version 8.2 (z/OS supports DB2).
- WebSphere Application Server, Version 5.1.1.
- Web server plug-ins.

Note: WebSphere Application Server V5.1.1 ships the following predefined applications that provide ARM-instrumentation support by default because they use the ARM instrumentation framework that is provided by WebSphere Application Server. They include but are not limited to the following:

- Apache HTTP Server
- Covalent Enterprise Ready Server
- Covalent Enterprise Ready Server FastStart
- Covalent FastStart Server
- IBM HTTP Server
- Lotus Domino® Enterprise Server
- Microsoft Internet Information Services (IIS)
- Sun ONE™ Web Server, Enterprise Edition

- IBM has ARM-instrumented the following non-WebSphere Application Server Web serving plug-ins which use the same filter types as the WebSphere Application Server Web serving plug-ins:
  - Independent plug-in for Apache HTTP Server 2.0.47 or later
  - Independent plug-in for IBM HTTP Server 2.0.47 or later
  - IBM HTTP Server for iSeries
- Non-IBM operating systems such as the following can be used as managed servers:
  - Windows 2000 Server Family
  - Windows 2003 Server Family
  - Solaris™ 8 and 9 supporting 32 and 64 bit user programs
- An application can be ARM-enabled should you want to use these facilities when the middlewares are not yet ARM instrumented.

Figure 3-2 describes the key elements of the EWLM structure and the general relationships between those elements.

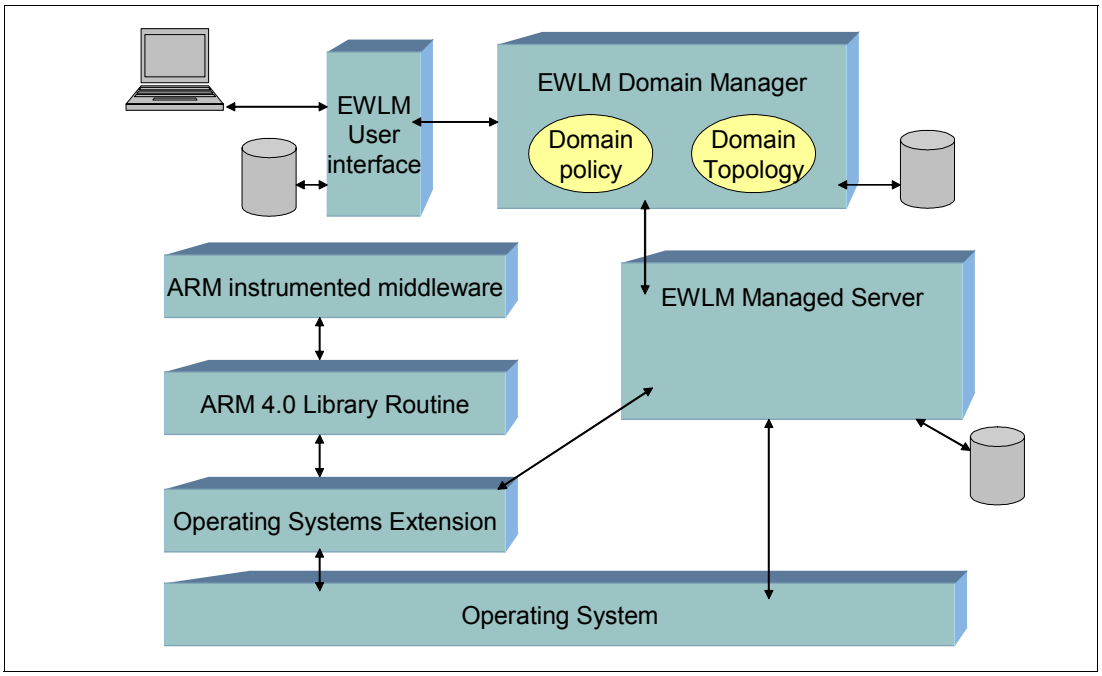

Figure 3-2 EWLM architecture

To learn more about EWLM, refer to the following:

- *IBM Enterprise Workload Manager*, SG24-6350
- *Hardening the EWLM Performance Data*, REDP-4018
- *Enterprise Workload Manager - Interpreting Control Center Performance Reports*, REDP-3963
- *Enterprise Workload Manager Overview*, REDP-3989
- *Enterprise Workload Manager z/OS Support*, REDP-3993

3.2 IBM Tivoli Provisioning Manager

Provisioning is a component of the Virtualization Engine that enables customers to share a pool of servers among different applications. These resources can be provisioned to the workload that has the biggest actual need and repurposed automatically to support a different workload if demand changes. The suite contains an integrated version of the Tivoli Provisioning Manager and platform provisioning functions for all IBM @server products with their major operating systems.

Figure 3-3 represents the main Tivoli Provisioning Manager components and their relationships.

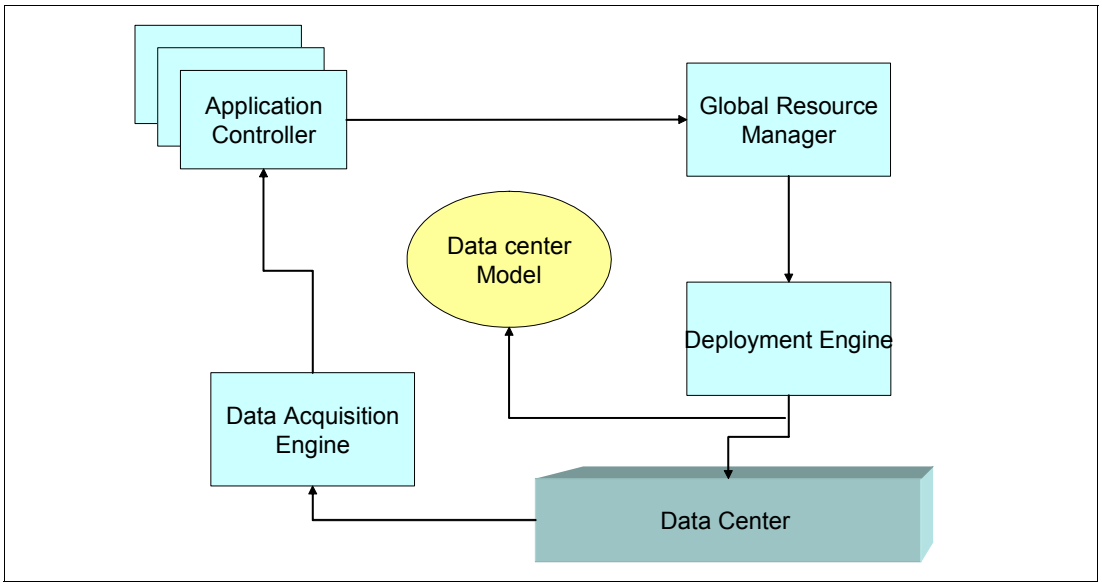

Figure 3-3 Tivoli Provisioning Manager components overview

3.2.1 The data center model

A data center is a complex environment. The data center model (DCM) provides a simplified view of resources by creating an abstraction layer that makes it easier to work with, allowing you to understand what resources you have and how they are being used to meet your business objectives. This component represents the physical and logical assets under the management of Tivoli Provisioning Manager, such as servers, switches, application software, VLANs, and so on. It keeps track of the data center hardware and associated allocations to customer sites.

The DCM attempts to model a physical data center as closely as possible. The objects that are defined within the DCM can best be described as being in two major categories:

- Infrastructure assets, which provide the fabric on which the various computing resources are built. These are both physical and logical.
  - The physical network assets include devices such as load balancers, switches, routers, firewalls, and so forth.
  - The logical assets include VLANs, IP subnets, and protocols.
- Application assets, which describe the hardware and software objects that provide the services defined by the data center.
  - The physical assets can include servers, clusters, and software products.
  - The logical assets can include resource pools, software stacks, customer SLAs, and so forth.

There are two ways to define resources into the DCM:

- Manually through the Tivoli Provisioning Manager operator user interface.
- Using an XML file and import it.
  a. Define data center resources as objects in an XML file (for example, define a server as a Server object in the XML file).
  b. Import the file, which the DCM uses to create a model of the data center.

Within the visual interface, most of the assets are defined through the following tabs: the first tab and the data center assets and resources tab. Some of the logical application assets are defined on a third tab, the customer applications tab. There are interdependencies among the objects. For example, before you can define a new virtual local area network (VLAN), you need to create a subnet. Also, to create and configure a resource pool or an application cluster, you must have a VLAN.

You must create and configure new assets and resources in Tivoli Provisioning Manager in the following order:

1. Switch fabric
2. Subnet
3. VLAN
4. Switch
5. Load balancer
6. Resource pools
7. Servers

Figure 3-4 shows the relationships between the objects that you define in your DCM.

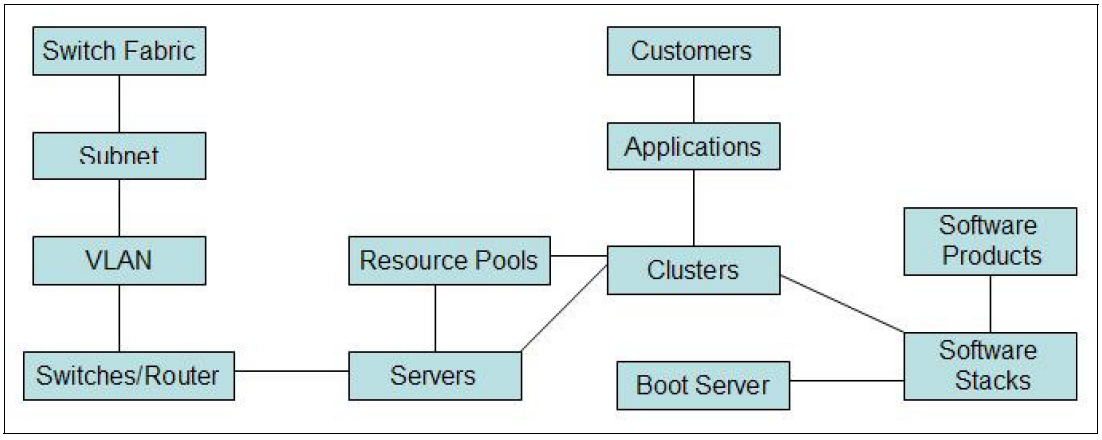

Figure 3-4 Relationships between objects in the DCM

The switch fabric has subnets, a subnet has a VLAN, and VLANs are used on switches, which might be extended to routers. Resource pools contain servers, which attach to switches. Resource pools, service clusters and servers can be dedicated to application clusters. Customers have applications and applications have clusters. Software products are contained within software stacks, which are assigned to clusters. A software stack might be contained on a boot server.

There are many DCM objects not referenced here. A more complete guide is available in *Provisioning On Demand: Introducing IBM Tivoli Intelligent ThinkDynamic Orchestrator*, SG24-8888.

You can display graphically the data center. To see a graphical view of the data center, select the overview link on the home page of the DCM Web interface. Figure 3-5 is an example of such a view. If you hover over any of the objects, a description of the various components within the data center appears. None of the objects can be manipulated in this view.

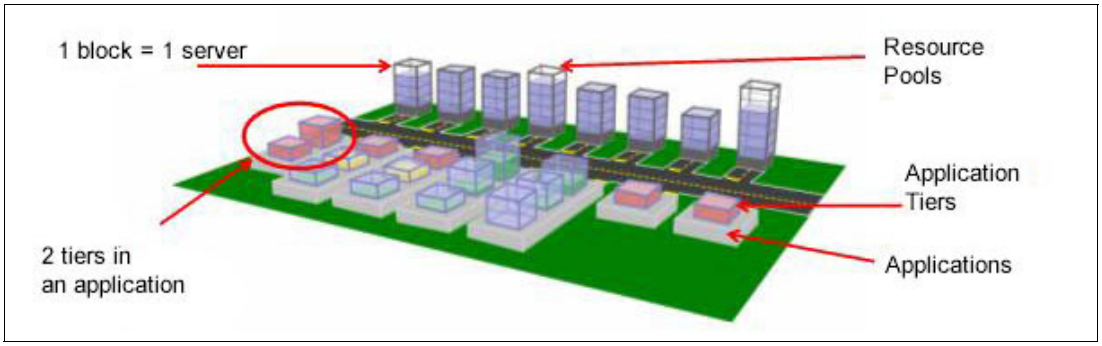

Figure 3-5 Graphical view of the data center

Tools to populate the DCM

The DCM can be built by using the graphical Web interface or by importing definitions contained in one or more XML files. Each of the two methods have their advantages and disadvantages as follows:

- The Web management GUI is simple to use but can be time consuming for repetitive tasks, such as defining multiple ports on a switch (for example 48 ports require 48 different definitions).
- Building the DCM using XML files is useful for repetitive tasks. However, creating the XML can be complex, time-consuming, and prone to errors. On the other hand, the XML can be reused, copied, or extended. Generally speaking, using the **`XMLImport`** command is the recommended method of updating the DCM.

A DCM cannot be reused on any other project using either method because there is no export facility in the current version. This means if you begin by using the XML method and then use the GUI method, your original XML does not reflect what is defined in the data center.

3.2.2 The workflow

A workflow is a list of steps or tasks that must be accomplished in order to achieve a desired goal. They also provide the ability to repeat that process in a consistent, auditable, and error-free manner.

In Tivoli Provisioning Manager terms, a workflow is a representation of the set of steps that must be followed in a data center environment in order to carry out a provisioning operation. The individual steps that make up a workflow are called *transitions*. Because workflows are used to carry out provisioning operations in a data center, the transitions enable a wide variety of activities to be performed on systems and network devices. Some examples of the types of operation that can be performed with workflows are:

- Allocate servers to clusters.
- Remove servers from clusters.
- Install software and patches.
- Uninstall software.
- Configure servers.
- Configure network devices.
- Run commands.

Workflows can be developed using XML and later imported into the Tivoli Provisioning Manager environment.

Workflows can be run in several different ways:

- A workflow that is assigned to an object can be run when the associated logical operation is triggered by:
  - The Resource Manager
  - SOAP commands
- Workflows not associated with logical operations can be executed from a GUI.

Workflows are controlled by the deployment engine.

3.2.3 The automation package

An automation package is collection of commands, shell scripts, workflows, logical operations, and Java plug-ins that applies to the operation of a specific type of software component or a physical device. Also referred to as *tc driver*.

3.2.4 The deployment engine

This workflow engine component of Tivoli Provisioning Manager runs repeatable workflows that automate the server configuration and allocation in the system. A workflow can represent either an entire reconfiguration process affecting multiple servers or a single step in a larger reconfiguration process.

3.2.5 For more information

To learn more about Tivoli Provisioning Manager, refer to the following:

- *An Introduction of IBM Tivoli Provisioning Manager Workflows*, TIPS0447
- *An Introduction to IBM Tivoli Provisioning Manager Automation*, TIPS0448
- *Developing Workflows and Automation Packages for IBM Tivoli Intelligent ThinkDynamic Orchestrator*, SG24-6057
- *Provisioning On Demand: Introducing IBM Tivoli Intelligent ThinkDynamic Orchestrator*, SG24-8888

3.3 IBM Director Multiplatform

IBM Director Multiplatform is the component of the Virtualization Engine that enables you to centrally manage a complex cross-platform environment that contains numerous servers, desktop computers workstations, mobile computers and assorted devices. It provides a single point of control for managing an heterogeneous environment. IBM Director Multiplatform is a comprehensive systems management solution, based on IBM Director, that can manage IBM iSeries, pSeries, xSeries, and BladeCenter servers, as well as desktop systems.

Note: IBM Director today is offered for IBM xSeries and BladeCenter servers only, providing tasks to exploit advanced management features. IBM Director Multiplatform is not a replacement for IBM Director. It allows IBM Director Server to run on all the platforms.

Important: In the future, it is expected that IBM Director Multiplatform and IBM Director the will merge. So, we use IBM Director as a generic name for both products, unless we refer to a specific function.

Also not that IBM plans are subject to change without notice and this information represents goals and objectives only. Such information is not intended as a definitive statement of a commitment to specific levels of function or delivery schedules.

IBM Director Multiplatform is comprised of the following components:

- IBM Director Server
- IBM Director Agent
- IBM Director Console

3.3.1 The IBM Director Server

The IBM Director Server is the main component — or aggregation point — for managing the agents. It contains the management data, the server engine, and the application logic. More than 44 functions are available through IBM Director Server. For example:

- The inventory function collects data on the hardware and software that is installed on a managed system in order to maintain control over the physical pieces of the environment.
- The Resource Monitoring function gathers statistics about critical system resources (for example microprocessor, disk, and memory usage) and sets thresholds to detect potential problems and alert users.
- Event Action Plans functions can define event filters and associated actions for resource and process monitors.

The IBM Director Server stores the management information in a database. You can access information that is stored in this integrated, centralized, relational database even when the managed systems are not available.

The IBM Director Server must be installed on each management server.

3.3.2 The IBM Director Agent

The IBM Director Agent gives you the ability to collect information from your heterogeneous environment and manage it from a single point of control. The IBM Director Server can communicate with all systems in the network on which the IBM Director Agent is installed. IBM Director Agent capabilities will also vary according to operating systems and hardware platforms. IBM Director Agent provides management data to the management server through various network protocols.

The IBM Director Agent must be installed on each management server and each managed server.

3.3.3 The IBM Director Console

The IBM Director Console is the user interface to the IBM Director Server. It is a Java-based console that is included with IBM Director. You can use the IBM Director Console to conduct comprehensive systems management using either drag-and-drop or single-click implementation. You can also use IBM Director Console to group managed objects, view associations and start tasks, as well as set IBM Director options and preferences. The IBM

Director Console can be integrated with the Virtualization Engine console for a consolidated view of systems management functions and control.

The IBM Director Console must be installed on the management server and on any system from which a system administrator will remotely access the management server (called a management console).

3.3.4 For more information

To learn more about IBM Director, refer to the following:

- *IBM @server* BladeCenter Systems Management with IBM Diector V4.1 and Remote Deployment Manager V4.1, REDP-3776
- *Monitoring Redundant Uninterruptible Power Supplies Using IBM Director*, REDP-3827
- *Implementing Systems Management Solutions using IBM Director*, SG24-61881

3.4 IBM Virtualization Engine console

The Virtualization Engine console is the Virtualization Engine user interface for administrators. It provides data to the browser, which allows a system administrator to connect to and view system health or launch other management consoles from any workstation client connected to the network without requiring a separate management console installation.

The Virtualization Engine console focuses on system health and resource monitoring and has the following benefits:

- Consolidates the view of the health of all system resources and provides a way to take corrective action.
- Provides one place to launch a variety of Virtualization Engine functions and management consoles.
- Provides a consistent user interface.

The goal of the Virtualization Engine console is to enable you to manage entire solutions — rather than specific products — by removing operating system boundaries and maximizing the sharing of resources. The Virtualization Engine console is planned to be the user interface identity of the Virtualization Engine and the interface into the systems services.

3.4.1 The Virtualization Engine console architecture components

The Virtualization Engine console is a Web-based interface built on the IBM Integrated Solutions Console, which is built on IBM WebSphere Portal technology as shown in Figure 3-6. The Integrated Solutions Console is a Web-based interface that provides a single platform that is capable of hosting all administrative console functions (setup, configuration, monitoring, and control) built by IBM server, software, and storage products.

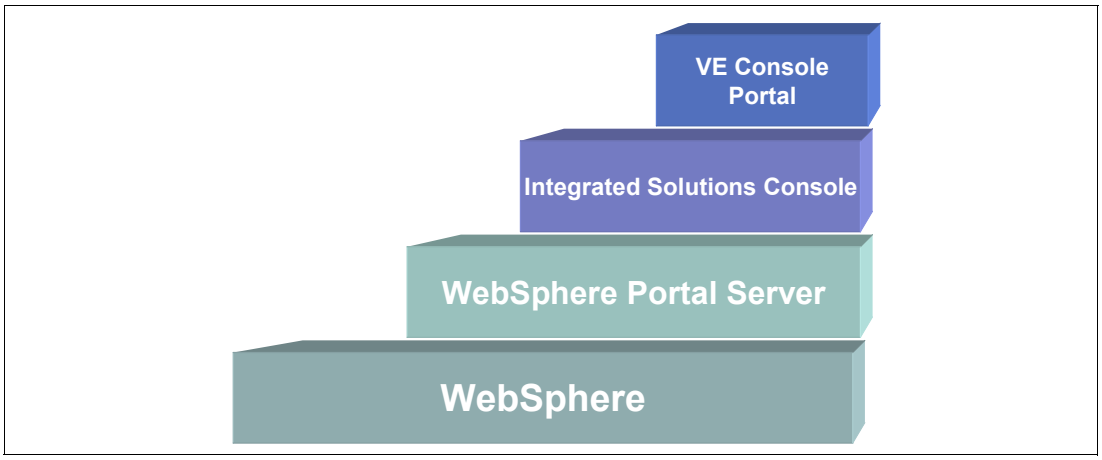

Figure 3-6 Positioning the Virtualization Engine console portal

The Virtualization Engine console needs to be installed on only one system. The Virtualization Engine console provides a single point of control, or *dashboard*, from which you can manage your heterogeneous IT system environment.

The basic architecture components

The Virtualization Engine console support is comprised of the following elements, as illustrated in Figure 3-7:

- The Web browser
- The Virtualization Engine console itself
- The Virtualization Engine console bridge
- The management source
- A set of managed resources

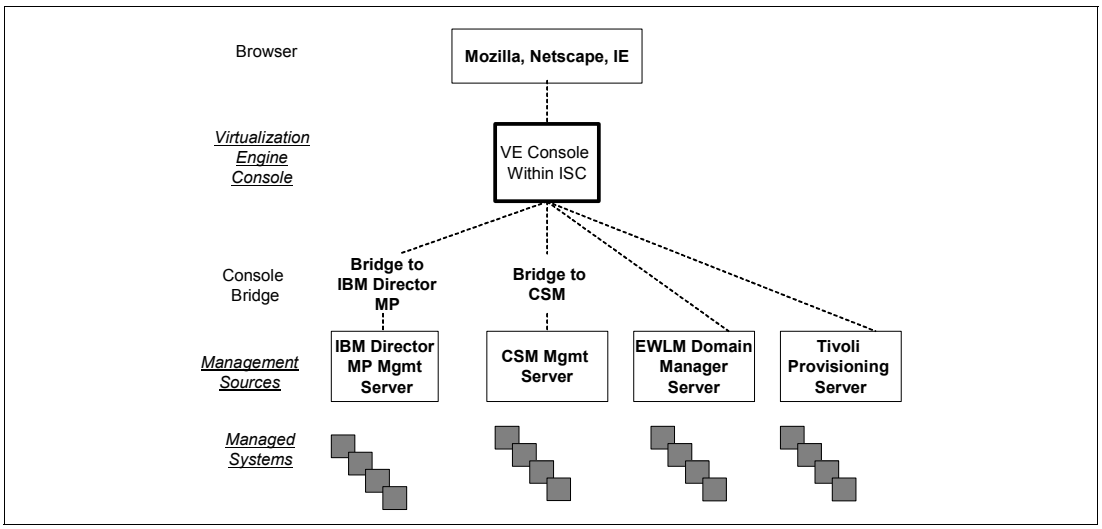

Figure 3-7 Virtualization Engine console architecture components

The following list describes these basic architecture components:

- The *browser* is used to access the Virtualization Engine console. Because the Virtualization Engine console is a Web-based application, you need to use a Web browser to access the console. Basically, the Virtualization Engine console works with any standard Web browser (such as Internet Explorer, Netscape, Mozilla, and so on) and does not require any additional plug-in. However, if you are planning to use full graphical features such as the dashboad, you need to install an SVG plug-in for your browser. The SVG support is currently available only on Internet Explorer 4.0 or higher.
- The Virtualization Engine *console bridge* is a software piece that allows the Virtualization Engine console to communicate via Web services with consoles. The Virtualization Engine console bridge provides interfaces to the Virtualization Engine console for accessing the Management Sources. The Virtualization Engine console bridge is running in WebSphere Application Server. Therefore, when you install the bridge, a WebSphere Application Server is installed with it. Usually, the Virtualization Engine console bridge must be installed on the same machine as the management source. However, if you are planning to use IBM Director, you can install it on a system other than that where IBM Director is installed.
- The *management sources* collect, distribute, and exchange data with the managed systems. The management source usually communicates with the Virtualization Engine console through a bridge. The management source is a previously deployed software server management application solution that the Virtualization Engine console is able to interact with, usually via a bridge. All information that you can see on the Virtualization Engine console are provided by the management sources.
- *Managed systems* are the endpoint systems, devices, or resources that are supported and managed by the management source, such as Cluster Systems Management, Management Central, Tivoli Monitoring, or IBM Director for Multiplatform). There is no Virtualization Engine unique code that is installed on the managed resources. However, managed agents typically run on the managed resources (for example, IBM Director Multiplatform Agent and EWLM Agent).

The Virtualization Engine console requires a WebSphere application server. In addition, Virtualization Engine console requires an LDAP server for its user management. The LDAP server is not installed with the Virtualization Engine console. Thus, you need to install it separately.

The Integrated Solutions Console

An Integrated Solutions Console instance consists of a framework and a set of one or more console components that are provided by product development groups (such as the Virtualization Engine console). Administrative activities are executed as portlets that have access to all of the standard services within the Java 2 Enterprise Edition (J2EE™) environment that is provided by the underlying the WebSphere application server.

The goal of the Integrated Solutions Console is to provide a single platform to host all administrative console functions. IBM plans to use the Integrated Solutions Console as the common user interface for management consoles. Customers and IBM Business Partners can also integrate their management consoles into the Integrated Solutions Console. The Integrated Solutions Console develops a single look and feel for common systems administration.

Portlets are built and integrated into packages that can then be deployed to the Integrated Solutions Console framework, which uses some WebSphere portal technology. The portlets can use the Integrated Solutions Console directly, they can use portal APIs, and they can use WebSphere Application Server APIs to get the job done. The administrators just access it from their browsers, and they do not have to have anything installed on their workstations.

In this first release, Integrated Solutions Console is using an IBM proprietary portlet API. However, in the future, the Integrated Solutions Console will provide the JSR168 API that will be available from the Java community process.

Figure 3-8 on page 49 illustrates the Integrated Solutions Console layout with the following working areas:

- A banner area, from where you can enter and customize user IDs, profiles, favorites, and so forth
- A navigation area, from where you can use the Virtualization Engine console
- A work area that is a Web application portlet controlled area

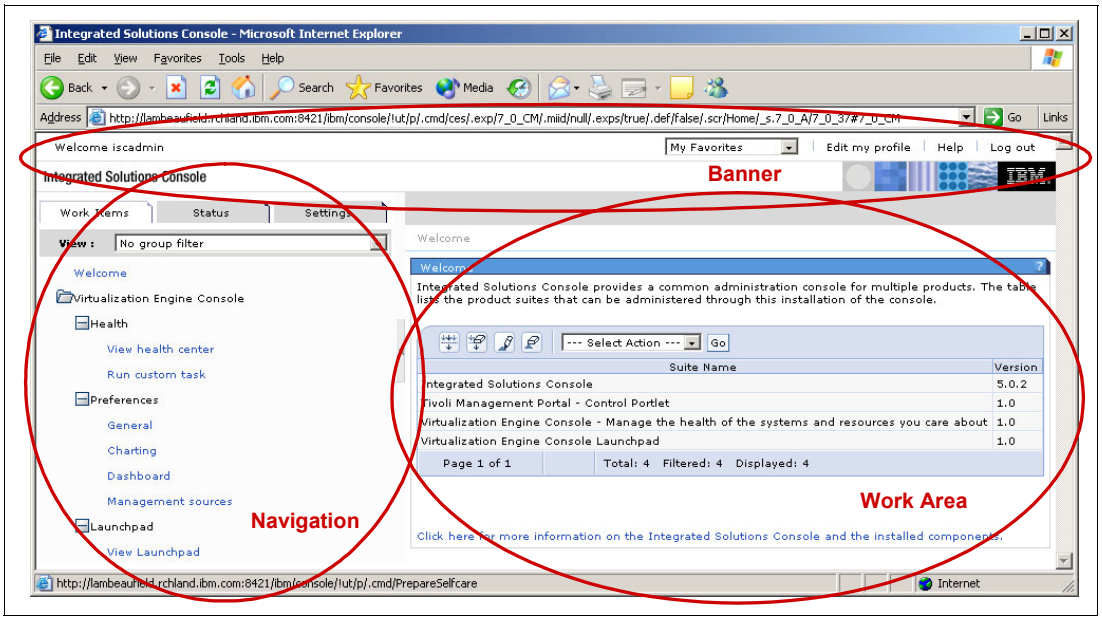

Figure 3-8 The Integrated Solutions Console layout

3.4.2 The Virtualization Engine console user components

The Virtualization Engine console consists of the following user components, as shown in Figure 3-9 on page 50:

- The *Launchpad* provides a single way to launch IBM @server system, application, and storage consoles.
- The *Health Center* provides a single user interface data to system status data that is captured from management sources.

The two tabs that are available are:

- The Settings tab provides the functions to customize and to administer the Virtualization Engine console.
- The Work Items tab provides access to the Launchpad or the Health Center.

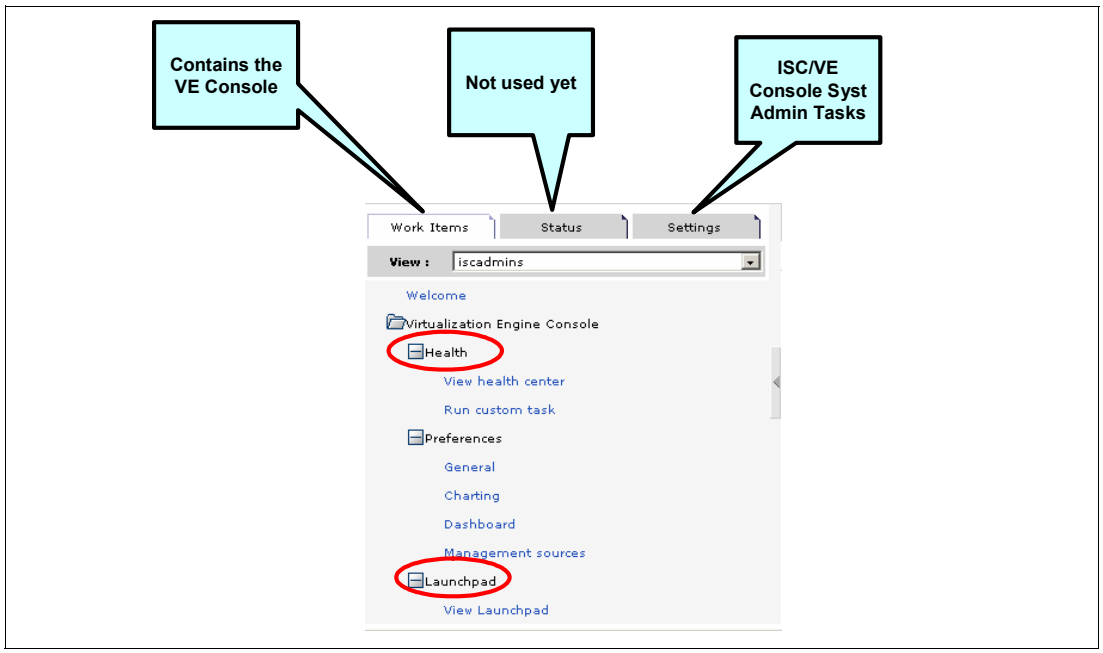

Figure 3-9 The Virtualization Engine console users components

The Launchpad

The Launchpad, as shown in Figure 3-10, provides a single view where you can launch your key IBM @server systems and storage consoles. The management console Launchpad provides links to installed and Web-based administrative consoles. These links launch administrative consoles that are part of the Virtualization Engine. Users can also customize the Launchpad by adding launch-points to custom consoles.

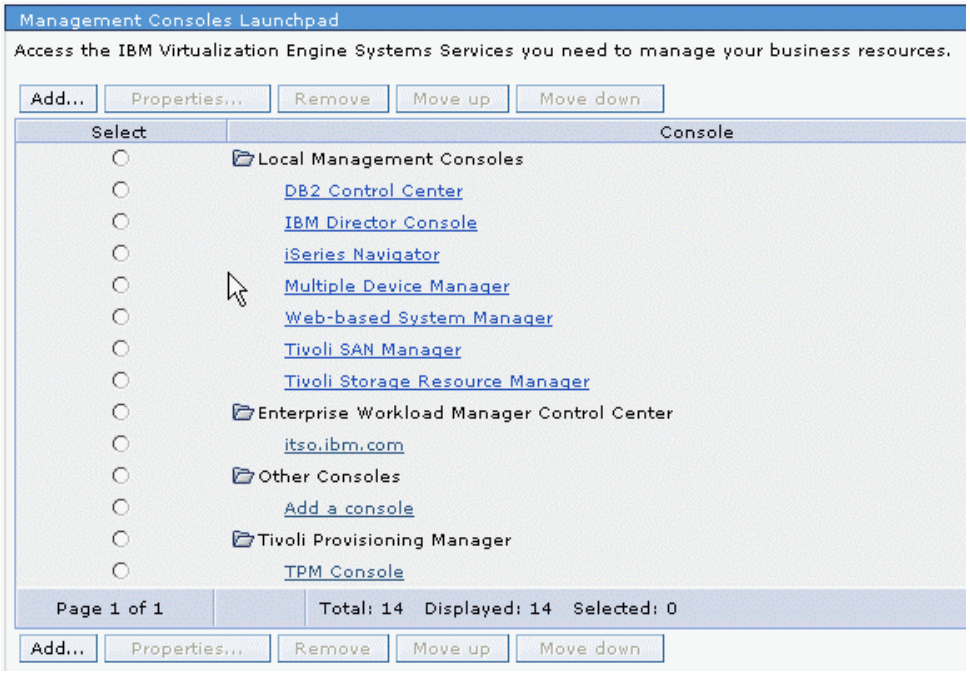

Figure 3-10 The Launchpad

You can think of the Launchpad as being similar to a folder of favorites or bookmarks that you use to reach your favorite Web sites. The first time you enter the Launchpad, there are three default folders:

- Local management consoles
- Enterprise Workload Manager Control Center
- Other consoles

The default local management consoles folder provides links to the local console software installed on your client system. That is the system where the browser is running, not the system where the Virtualization Engine server is installed. In other words, if you want to launch the IBM Director Console from the Launchpad folder using a browser on your ThinkPad, then you must have the software for the IBM Director Console installed on your ThinkPad.

Because links to local consoles are enabled only if they are installed on the user's browser system, some local management consoles will not or cannot be run on some operating systems. For example, iSeries Navigator and Tivoli SAN Manager do not run on Linux. Therefore, these consoles could not be launched from a browser running the Virtualization Engine console on those platforms.

Local management consoles cannot be added, removed, or modified. The currently supported default local management consoles are:

- DB2 Control Center: manage systems, DB2 UDB instances, databases, and database objects such as tables and views.
- IBM Director Console: user interface for the IBM Director which is a powerful systems management source.
- Multiple Device Manager: reduces the complexity of SAN management by allowing administrators to configure, manage, and monitor performance of storage from one console.
- iSeries Navigator: GUI for managing and administering iSeries servers from Windows desktops.
- Web-based System Manager: a comprehensive set of system management tools for AIX 5L which allows access to and management of AIX systems from a remote GUI.
- Tivoli SAN Manager: comprehensive management interface for multivendor SANs.
- Tivoli Storage Resource Manager: comprehensive capacity management tool for heterogeneous storage environments. It includes reporting, monitoring and management of direct attached storage, network attached storage, and SANs.

Launchpad can launch client-based management consoles (such as IBM Director Multiplatform or Cluster Systems Manager) through the use of a Virtualization Engine console bridge. These consoles are part of the Integrated Solutions Console — in other words, they appear in a portlet on the Virtualization Engine console.

Launchpad can also launch browser-based management consoles which are not yet part of the Integrated Solutions Console (such as those for Tivoli Provisioning Manager and Enterprise Workload Manager).

Launchpad provides a single starting point for your management consoles. The Launchpad can call the Virtualization Engine console components as well as local consoles that are not part of Virtualization Engine.

Launchpad is a list of links to consoles. You launch each console by clicking a link. If you are planning to add a Web-based management console in the list, you can add it by providing the

URL and port information. However, if you are planning to call local management consoles, it is not possible to add another console which is not listed. In addition, in order to start local consoles, a Java Web start plug-in is required for your browser, and each local console must be installed on the machine that executes the Web browser.

The Health Center

The Health Center, as shown in Figure 3-11, provides a single user interface in a heterogeneous environment by aggregating information where you can perform key server (such as AIX, Linux, i5/OS, and Windows) and storage monitoring and management tasks. The Health Center simplifies the work of a system administrator by providing a Web console that displays the health of the server and storage resources, monitors, and tasks. You can identify problems and take actions to solve them from the Health Center by clicking the resource, monitor, or task that needs attention. The Health Center provides a bird's eye view of system resource health across a converged system and storage management environment.

The Health Center provides the following functions:

- Displays and analyzes the status and health of resources.
- Displays real-time resource data and analyses through the dashboard display.
- Monitors resources across a distributed cross-platform environment.
- Runs actions against resources in response to conditions detected by the Health Center.

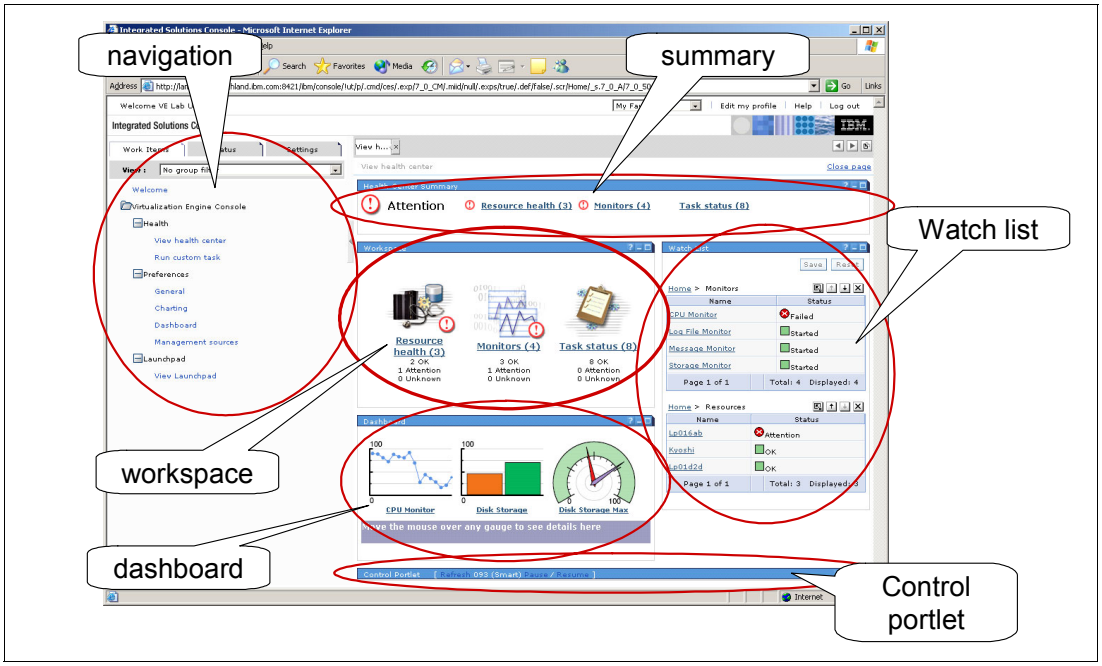

Figure 3-11 The Health Center

The Health Center gets its underlying information from management sources and consolidates health from a variety of management sources. The Health Center is a portlet-based console. It is a *portal* which comprises a set of related *portlets*. Portlets are little pieces of the user interface which are customizable. A portlet is a Web-based interface similar to a client-based dialog box or window. A set of portlets is a portal. A portal is just a single Web page broken into multiple portlets or sections. A user can work with multiple portlets at a time, all available and running in the Health Center.

As the name Health Center implies, it gives you a way to check the pulse and other vital signs of the managed systems in your IT infrastructure using a set of pre-installed management sources such as:

- IBM Director
- Multiple Device Manager (MDM)
- Cluster Systems Manager (CSM)
- Management Central (MC) for iSeries
- Tivoli Monitoring

The Health Center checks, displays, and analyzes the status and health of resources. It displays a resource's real-time data, analyzes it using charts and graphs, and adds it to a dashboard display. It monitors resources across a distributed cross-platform environment.

The main Health Center functions are as follows:

- The Resource Health
- The Monitors
- The Tasks

Figure 3-12 shows how these functions appear in the work space and the summary areas.

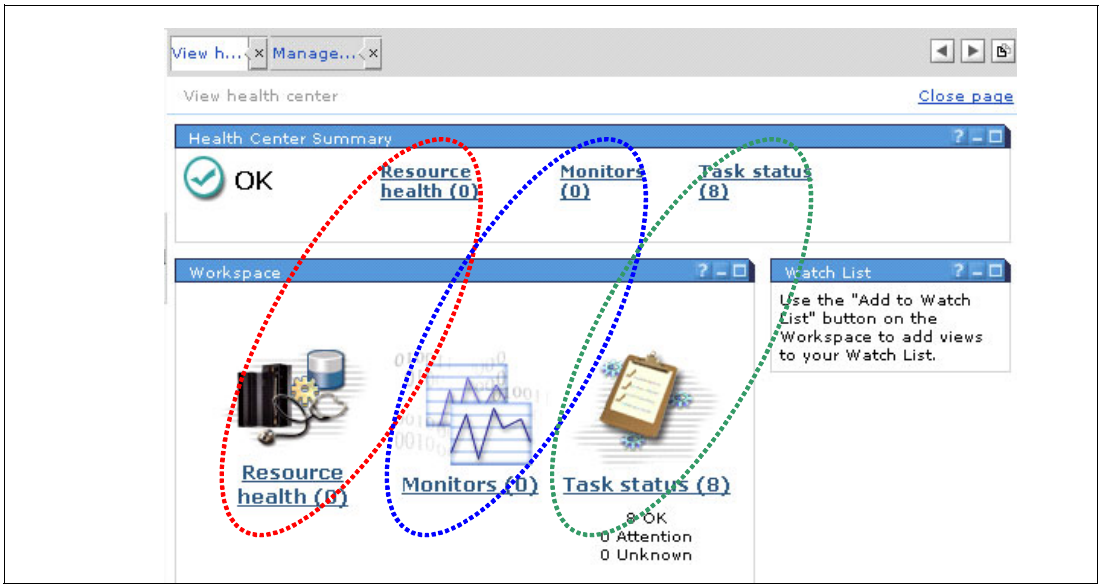

Figure 3-12 Heath Center functions

Resource health

Resource health lets you look at managed systems, such as servers, devices, and groups of systems. You can quickly navigate through several levels of resources and take action on a resource that might need fixing. When you open the Resource Health view, you see a list of systems and groups with status icons. You configure this list of systems in your preferences.

The underlying data in all cases flows from the managed system (such as a IBM Director Agent running on a server) to the management source (such as IBM Director Multiplatform), through a Virtualization Engine console bridge to the Health Center in the Virtualization Engine console. For example, you could keep track of processor utilization, events, logs, or processes on any of the servers managed by multiple IBM Director Multiplatform management sources. Resources include components such as BladeCenter. The Virtualization Engine console provides status and health information for each resource.

If you open the Resource Health window and click **Add Resources**, management source reports the list of resources. You can add or remove resources which are configured on management sources.

Figure 3-13 shows a part of the resource list that is reported by IBM Director server. If you want to add a resource, select that resource, and then click **OK**.

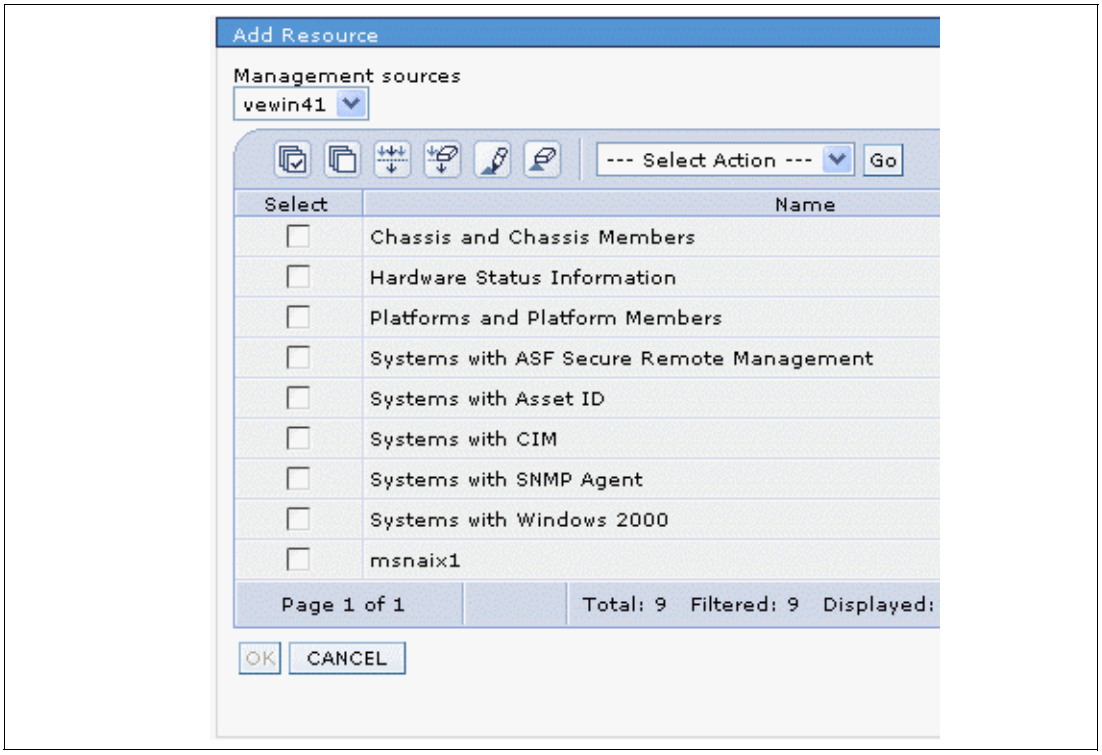

Figure 3-13 Example of a resource list that is reported by the IBM Director Server

If there is a problem on a resource, the Virtualization Engine console displays an `Attention` message, as shown in Figure 3-14.

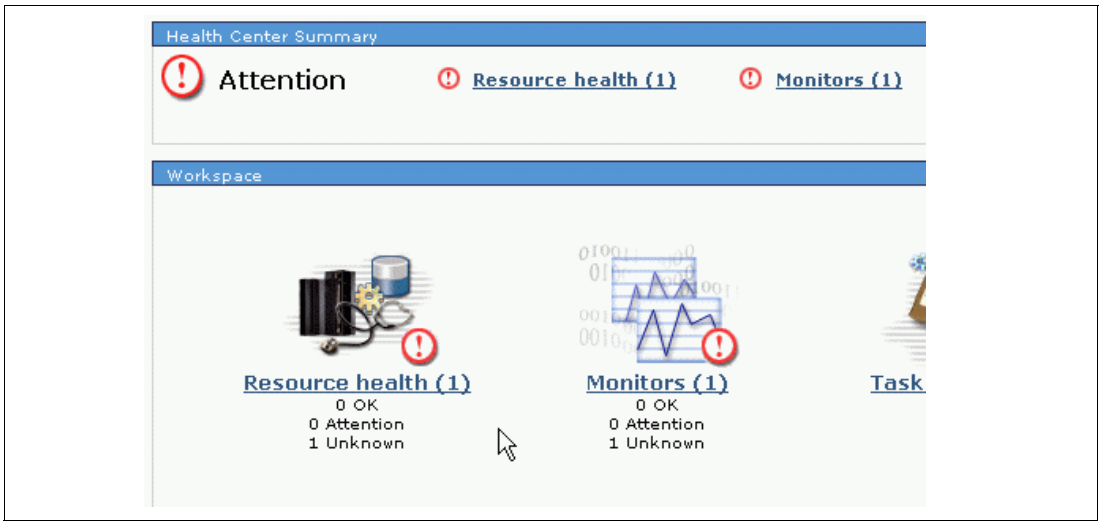

Figure 3-14 Example of a warning state that is provided by the Virtualization Engine console

Monitors

Monitors let you set work with monitors that you have set up on a management source (such as IBM Director Multiplatform). You then are notified when a critical event is triggered, and you can view a graphical representation of the managed systems. There are several ways to make the data from Monitors easy to view and organize:

- Real-time graphs and charts.
- Dashboard.

  The Dashboard portlet in Health Center displays a near real-time graphical representation of your managed systems, as shown in Figure 3-15. The metrics can be displayed in a variety of formats, so you can pay attention to whatever you are most focused on at the moment. Of course, it is something you can customize to meet your needs. You can add and remove gauges and change their formats at will.

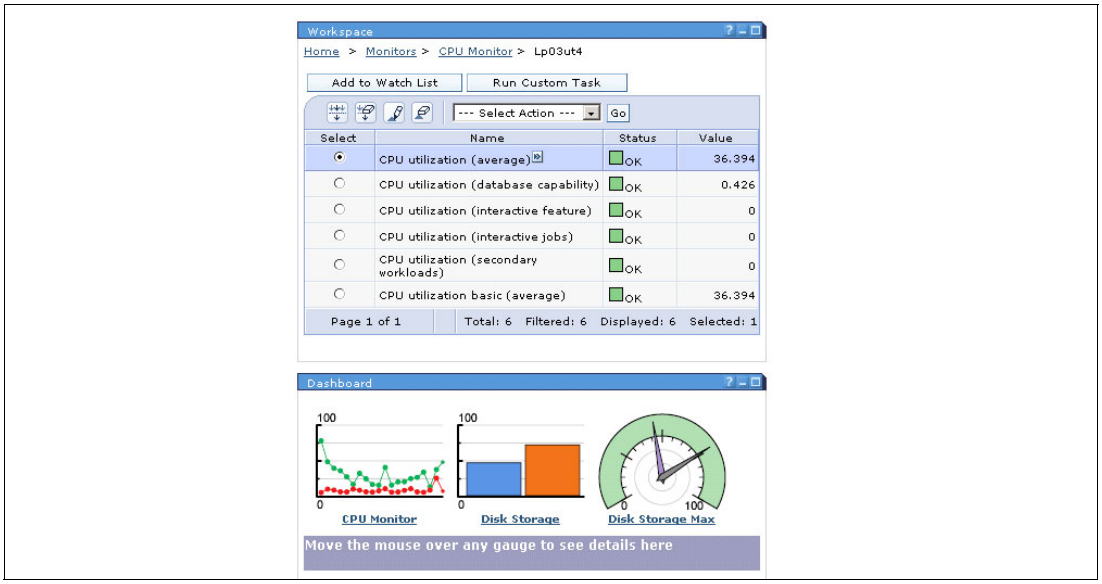

Figure 3-15 The dashboard

- Watch List.

  The Watch List portlet in the Health Center displays views of resources, monitors, and tasks. You can add any data items that you want to monitor closely to the Watch List. The information in the Watch List persists even as you navigate away from the Health Center view, change to another page in the Virtualization Engine console or log off the Virtualization Engine console. The Watch List, as shown in Figure 3-16 on page 56, enables you to watch selected views of resource health and navigate through the resources to watch a subset of data items contained in that resource view.

Figure 3-16 The Watch List

The kind of resources that can be monitored depends on the managed source. In order to add monitors, you need to configure management source to provide this information. For example, you need to configure Threshold in IBM Director management source to monitor it from the Virtualization Engine console. Figure 3-17 shows the monitors. If you want to add monitors, click **Add Monitor**, and the list of monitors that are configured on management source are displayed. Select the monitor that you want to add, and click **OK**.

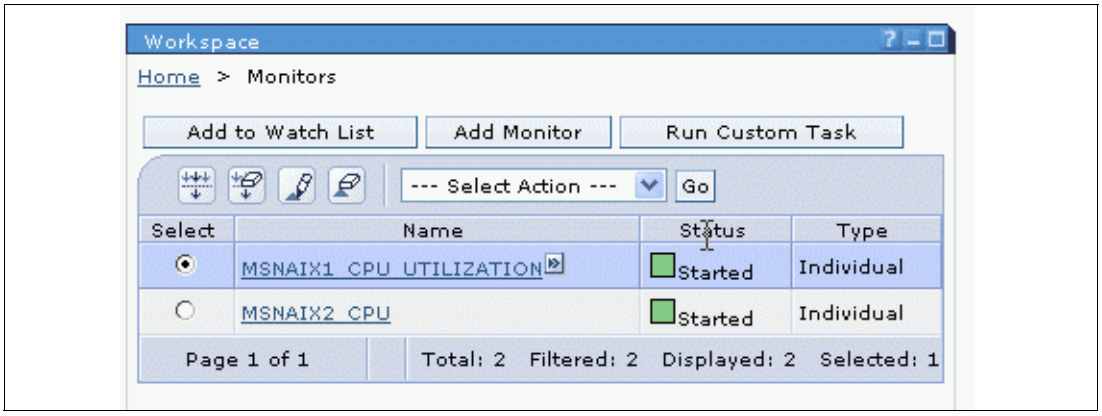

Figure 3-17 Example of systems that are monitored with the Virtualization Engine console

It is possible to display detailed data-like resources on Resource Health. In addition, you can display a graph for monitors. We need to install an SVG plug-in our Web browser to display a graph. Click the monitor name and select **Display graph**.

Figure 3-18 and Figure 3-19 show a sample graph of processor use.

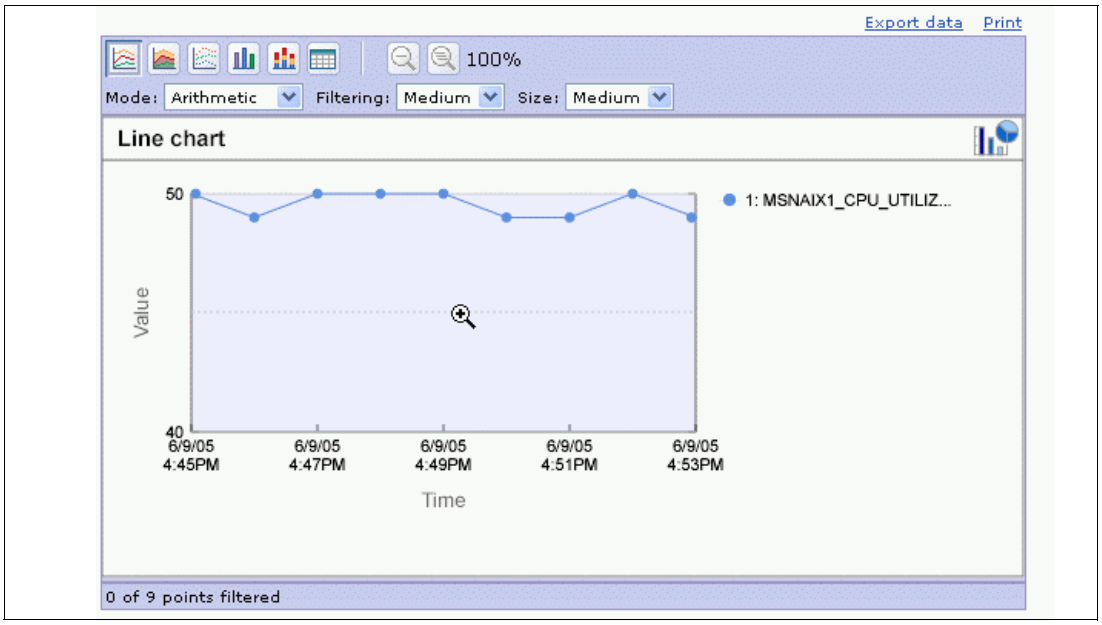

Figure 3-18 Example of display graph

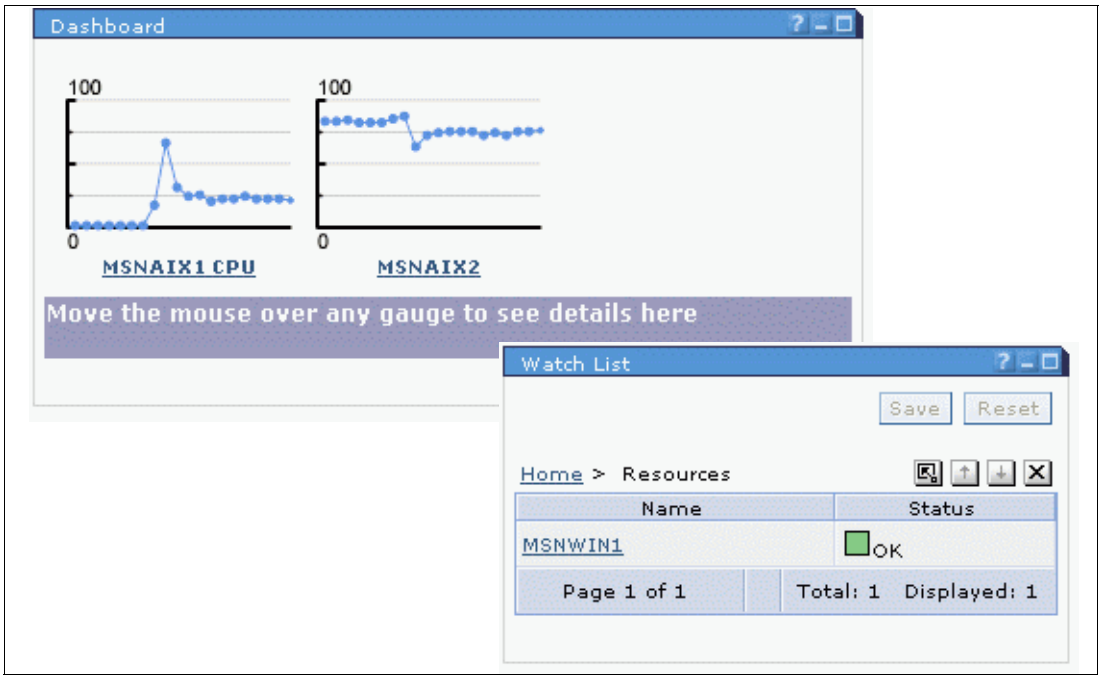

Figure 3-19 Example of a graph of monitor

We can configure dashboad and Watch List for a specified device. Dashboad provides a rough image to monitor resources visually. In order to use dashboad, you need to install SVG plug-in for our browser. On the other hand, the Watch List is a shortcut for any resource and monitor. When we register a resource or a monitor in the watchlist, we can check each status from the top page. Figure 3-20 on page 58 shows dashboard and Watch List.

If you want to add a resource or monitor in Watch List, click **Add Watchlist**. If you want to add a monitor in dashboad, click the monitor name. Then, select **Add to Dashboad**, as shown in Figure 3-20.

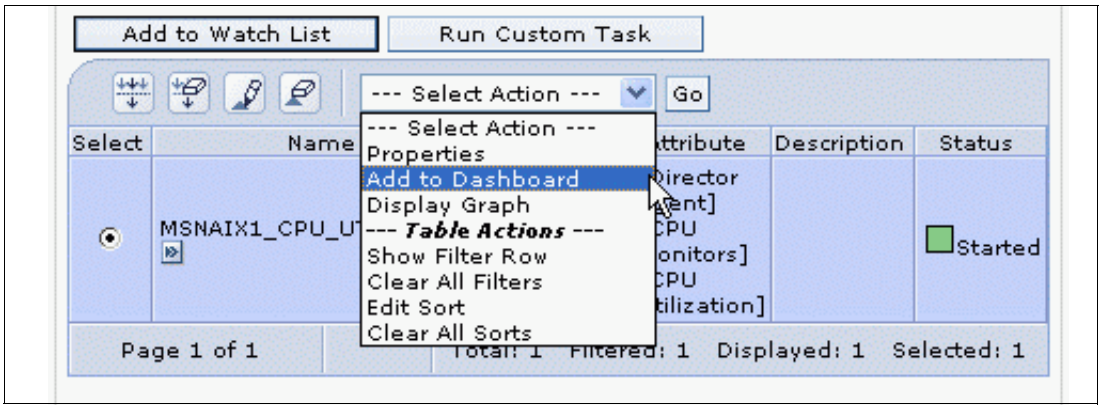

Figure 3-20 Example of adding a component on the dashboard

Make sure that the kind of data, information and task depends on which Management Source you use. Visualization Engine console is a consolidated interface for Management Sources and does not have functions to contact each managed resources directly.

Task Status

Task Status lets you run simple or customized tasks and check their status. The Health Center lets you perform simple actions such as starting and stopping resources and monitors. However, you can also use custom tasks to run more complex actions against your resources and monitors. Custom tasks are predefined on the management source. To run a custom task from the Health Center, you need to configure that task on the management source first (for example, IBM Director Multiplatform or iSeries Navigator).

The Virtualization Engine console Health Center supports IBM Director.

3.5 The common runtime

The common runtime of the Virtualization Engine is made of the pieces common to the different Virtualization Engine components. For example:

- An LDAP server is needed for the Virtualization Engine console and for Tivoli Provisioning Manager
- An LDAP server requires a database
- The IBM Director and Tivoli Provisioning Manager require a database to store specific information
- Different administration interfaces for the Virtualization Engine console and Tivoli Provisioning Manager use Java programs which need an application server to be executed
- The Virtualization Engine promotes Tivoli Directory Server as the LDAP server, DB2 as the database, and WebSphere as the application server (although some other configurations might be possible as written in this book when needed)

The idea of the common runtime is to use the same LDAP server, the same database, and the same application server for all the installed Virtualization Engine components. In the Virtualization Engine Version 1, the common components can be still requested at different

levels, requiring multiple versions or releases of the same product. The intent over time is to deliver one common set of these common components, making the installation smarter and smoothly.

3.5.1 The components

The Virtualization Engine common runtime consists of the following components:

- Integrated Solutions Console Version 5.0.2
- WebSphere Application Server Version 5.1
- DB2 Universal Database Server Version 8.1
- IBM Tivoli Directory Server Version 5.1 and 5.2

Integrated Solutions Console 5.0.2

The Integrated Solutions Console constitutes the Web-based administration interface for the Virtualization Engine console. This is the central management portal for managing a Virtualization Engine environment. The Integrated Solutions Console provides an interface for managing each Virtualization Engine system services component either directly from the Web-based GUI or by providing a launchpad to a particular components from the management console.

WebSphere Application Server 5.1

This is a core component of Virtualization Engine that is used by the following system services components:

- Virtualization Engine console
- EWLM Domain manager
- Tivoli Provisioning Manager

WebSphere provides the Web-based application interface for managing and maintaining these Virtualization Engine system service components.

DB2 Universal Database Server 8.1

DB2 is used by IBM Tivoli Directory Server to store the LDAP user and group account information. This information is used in application authentication by both the Virtualization Engine console and Tivoli Provisioning Manager. Additionally, DB2 is used to store IBM Director and Tivoli Provisioning manager data.

IBM Tivoli Directory Server

Tivoli Directory Server provides LDAP authentication for both the Virtualization Engine console and Tivoli Provisioning Manager system services components. These components require separate version of Tivoli Directory Server. In both cases, Tivoli Directory Server is installed as a prerequisite prior to the actual components installation:

- Tivoli Directory Server Version 5.1, used by Virtualization Engine console, installed using the Tivoli Directory Server installation wizard
- Tivoli Directory Server Version 5.2, used by Tivoli Provisioning Manager, installed using the Tivoli Provisioning Manager prerequisite installer

Important: Tivoli Directory Server 5.1 and Tivoli Directory Server 5.2 cannot be installed in the same operating system image.

3.5.2 Installing the common runtime

The common runtime components are installed by default through one of the following methods:

- Virtualization Engine installation wizard (Figure 3-21)
- Component specific installation wizard

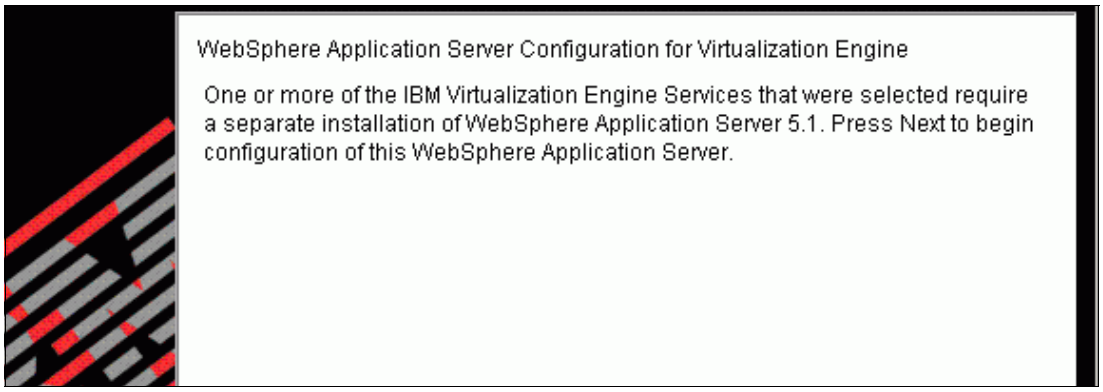
WebSphere Application Server Configuration for Virtualization Engine
One or more of the IBM Virtualization Engine Services that were selected require a separate installation of WebSphere Application Server 5.1. Press Next to begin configuration of this WebSphere Application Server.

Figure 3-21 WebSphere installation as part of the Virtualization Engine installation wizard

In most cases, you do not have to worry about the common runtime components. However, a circumstance where you do have to worry about the common runtime components is when you need to patch the particular component. This is the case with the IBM Director Multiplatorm requirement for DB2 8.1 Fix Pack 5, and the Tivoli Provisioning Manager 2.1 requirement for WebSphere 5.1, cumulative fix 3.

Another thing to understand is that some parts of the common runtime are shared by different Virtualization Engine system service components.

As an example, IBM Tivoli Directory Server 5.1 is a prerequisite of the Virtualization Engine console. It is installed via its own installation wizard. Tivoli Directory Server also installs DB2 8.1 as a database for storing its LDAP directory information. In a scenario where IBM Director is part of the environment, IBM Director can leverage this existing instance of DB2.

Note: IBM Director has several database options. In this book, we use DB2 because it is part of the Virtualization Engine common runtime.

In addition to the common runtime, Tivoli Directory Server and the Integrated Solutions Console contain their own embedded WebSphere Express instances. These are installed as part of each components own internal installation routine.

Figure 3-22 details the relationship between the Virtualization Engine system services components and the common runtime.

System Service	WebSphere Application Server 5.1	DB2 8.1	IBM Tivoli Directory Server	Integrated Solutions Console 5.0.2
EWLM Domain Manager	Yes	N/A	N/A	N/A
IBM Director Multiplatform	N/A	Yes, separate installation	N/A	N/A
Virtualization Engine Console	Yes	N/A	Yes, version 5.1	Yes
IBM Tivoli Provisioning Manager	Yes	Yes	Yes, version 5.2	N/A
IBM Grid Toolbox	N/A	N/A	N/A	N/A

Installed as part of the Virtualization Engine installation wizard

Installed as part of component specific installation wizard

Figure 3-22 Common runtime components

3.6 IBM Grid Toolbox V3 for Multiplatforms

The grid computing contribution is to enable and simplify collaboration among a very large audience, while offering important standards that enable very heterogeneous systems to work together to form the image of a large virtual computing system offering a variety of virtual resources. The goal to virtualize the resources on the grid and more uniformly handle heterogeneous systems creates new opportunities to better manage a larger, more disperse IT infrastructure. A grid federates a large number of resources contributed by individual machines into a greater total virtual resource. For applications that are grid-enabled, the grid can offer a resource balancing effect by scheduling grid jobs on machines with low utilization.

The IBM Grid Toolbox V3 for Multiplatforms implements the Open Grid Services Infrastructure (OGSI) standards and provides the tools to build a grid and to develop, deploy, and manage grid services. The IBM Grid Toolbox consists of the following:

- A hosting environment that is capable of running grid services and sharing them with other grid participants, such as grid service providers and grid service consumers.
- A set of tools to manage and administer grid services and the grid hosting environment, including a Web-based interface, the Grid Services Manager.
- A set of APIs and development tools to create and deploy new grid services and grid applications.

Restriction: This component is out of the scope of this redbook.

To learn more about Grid, refer to the following:

- *Grid Computing with the IBM Grid Toolbox*, SG24-6332
- *Grid Services Programming and Application Enablement*, SG24-6100
- *Introduction to Grid Computing with Globus*, SG24-6895
- *Fundamentals of Grid Computing*, REDP-3613

3.7 Resource Dependency Service

Resource Dependency Service (RDS) is a new offering part of the Virtualization Engine management collection, to help companies gain a view into IT resources. RDS can be used to help define and discover relationships between users, resources, applications, and networks. This allows:

- Automatic discovery of resources.
- The ability to map business processes to IT resources.
- Identifying dependencies and relationships
- Interoperability with Tivoli Change and Configuration Management database.
- A graphical display of topology via the Virtualization Engine console.

Restriction: This component is out of the scope of this redbook.

3.8 Topology and terminology

With the Virtualization Engine, we often refer to the following definitions:

- Virtualization Engine Management Server

  This is a platform with intelligence functions to issue instructions that enforce policies. The Virtualization Engine Management Server is a collector of monitoring data. Examples include EWLM Domain Manager, IBM Director Server, Systems Provisioning Management Server.

- Virtualization Engine Managed Node (or managed server)

  This is the platform that is managed or provisioned. It has intelligence to perform management functions on a single server basis. How a node is managed is the responsibility of the manager software. Examples include EWLM Managed Server and Systems Provisioning Provisioned Platform.

- Virtualization Engine common runtime

  These are technologies that provide functions which are leveraged by the Base or Priced Features for basic fabric capabilities. The Virtualization Engine is dependent on the presence of these technologies. Examples include WebSphere Application Server, Integrated Solutions Console, IBM Directory Server, and DB2.

- Managed from

  The system with the Virtualization Engine common runtime and components installed.

- Managed through

  The management server that is being used as the resource instrumentation for Virtualization Engine services.

- Managed to

  The end-point that is being managed.

3.9 Supported operating systems

You can find latest information about the supported operating systems at:

http://publib.boulder.ibm.com/infocenter/eserver/v1r1/en_US/index.htm?info/icmain.htm

The following tables summarize the supported operating systems and platforms at the time of the writing of this redbook. Table 3-1 describes the support on the zSeries.

Table 3-1 zSeries support

EWLM	Domain Manager	not supported
	Managed Server	z/OS V1R6
Tivoli Provisioning Manager 2.1	Managed through server	z/VM 4.4 (through VM API)
	Provisioned Server	SuSE LINUX Enterprise Server 8 with service pack 2
IBM Director Multiplatform	Console	not supported
	Server	not supported
	Agent	not supported
Virtualization Engine console	Console	not supported
	Bridge	not supported

Table 3-2 describes the support on the pSeries.

Table 3-2 pSeries support

EWLM	Domain Manager	AIX 5L 5.2 and AIX 5L 5.3
	Managed Server	AIX 5L 5.2 and AIX 5L 5.3
Tivoli Provisioning Manager	Management Server	AIX 5L 5.2 and AIX 5L 5.3
	Managed through server	► AIX 5L 5.2 and AIX 5L 5.3 (through NIM) ► Red Hat Enterprise LINUX AS 3.3 (through CSM) ► SuSE LINUX Enterprise Server 8 and 9 (through CSM)
	Provisioned server	► AIX 5.1, 5.2 and 5.3 ► Red Hat Enterprise LINUX AS 3.3 ► SuSE LINUX Enterprise Server 8 and 9
IBM Director Multiplatform	Console	not supported
	Server	not supported
	Agent	► AIX 5L 5.2 and 5.3 ► Red Hat Enterprise LINUX 3.0 and 3.3 ► SuSE LINUX Enterprise Server 8 and 9
Virtualization Engine console	Console	AIX 5L 5.2 and 5.3
	Bridge	AIX 5L 5.2 and 5.3 (through CSM)

Table 3-3 describes the support on the iSeries.

Table 3-3 iSeries support

EWLM	Domain Manager	► AIX 5L 5.2 and AIX 5L 5.3 ► i5/OS 5.3 ► Windows 2000 and Windows 2003
	Managed Server	► AIX 5L 5.2 and AIX 5L 5.3 ► i5/OS 5.3 ► Windows 2000 and Windows 2003
Tivoli Provisioning Manager	Management Server	► AIX 5L 5.2 and AIX 5L 5.3 ► SuSE LINUX Enterprise Server 8 ► Windows 2000 and Windows 2003
	Managed through server	OS/400 5.2 and i5/OS 5.3
	Provisioned server	► SuSE LINUX Enterprise Server 8 ► Windows 2000 and Windows 2003
IBM Director Multiplatform	Console	Windows 2000 and Windows 2003
	Server	Windows 2000 and Windows 2003
	Agent	► Windows 2000 and Windows 2003 ► i5/OS 5.3 ► Red Hat Enterprise LINUX 3.0 and 3.3 ► SuSE LINUX Enterprise Server 8 and 9
Virtualization Engine console	Console	► AIX 5L 5.2 and AIX 5L 5.3 ► i5/OS 5.3 ► Windows 2000
	Bridge	► AIX 5L 5.2 and AIX 5L 5.3 ► i5/OS 5.3 ► Windows 2000 and Windows 2003

Table 3-4 describes the support on the xSeries.

Table 3-4 xSeries support

EWLM	Domain Manager	► Red Hat Enterprise LINUX AS 2.1 and 3.2 ► SuSE LINUX Enterprise Server 8
	Managed Server	► Windows 2000 and Windows 2003 ► VMWARE ESX 2.5
Tivoli Provisioning Manager	Management Server	► Red Hat Enterprise Linux AS 2.1 and 3.2 ► SuSE LINUX Enterprise Server 8 Windows 2000 and Windows 2003
	Managed through server	► Red Hat Enterprise LINUX AS 3.0 (through CSM) ► Red Hat Enterprise LINUX AS 3.3 (through RDM) ► SuSE LINUX Enterprise Server 8 (through CSM or RDM) ► SuSE LINUX Enterprise Server 9 (through CSM) ► Windows 2000 and Windows 2003 (through RDM)
	Provisioned server	► Red Hat Enterprise Linux AS 2.1, 3.0 or 3.3 ► SuSE LINUX Enterprise Server 8 and 9 ► Windows 2000 and Windows 2003
IBM Director Multiplatform	Console	► Red Hat AS 2.1 and AS 3.0 ► SuSE LINUX Enterprise Server 8 and 9 ► Windows 2000 and Windows 2003 and Windows XP
	Server	► Red Hat AS 2.1 and AS 3.0 ► SuSE LINUX Enterprise Server 8 and 9 ► Windows 2000 and Windows 2003
	Agent	► Red Hat 2.1 AS ► ed Hat Enterprise Linux 2.1 and 3.0 ► SuSE LINUX Enterprise Server 8 and 9 ► Windows 2000 and Windows 2003 and Windows XP and Windows NT® ► VMware ESX 1.5.2 , 2.0, 2.0.1, 2.1 and 2.5 ► NetWare 6.0 and 6.5

Virtualization Engine console	Console	► Red Hat Enterprise LINUX AS 2.1 ► SuSE LINUX Enterprise Server 8 ► Windows 2000
	Bridge	► Red Hat Enterprise LINUX AS 2.1 and 3.0 (with IBM Director Server) ► SuSE LINUX Enterprise Server 8 ► Windows 2000 and Windows 2003 (with IBM Director Server)

4

The planning advisor tool

To help you plan your installation of the IBM Virtualization Engine platform, IBM developed the planning advisor tool. This interactive tool produces a customized planning output that is based on your input. The planning advisor does the following:

- Helps you to identify what systems services you want to install.
- Guides you through the design of your Virtualization Engine topology.
- Helps you obtain the information you need to order the Virtualization Engine.
- Lists the prerequisites that you need before installing the components.
- Lets you generate customized installation flow charts.
- Allows you to print and record your customized output before beginning the installation process.

Based on your input, the planning advisor generates recommendations to help you:

- *Decide* which systems services to include in your Virtualization Engine management collection.
- *Recommend* how to deploy the components of the Virtualization Engine across your enterprise topology.
- *Generate* customized documentation to help you prepare your environment and to help you install and configure the Virtualization Engine in your enterprise.

The planning advisor has been largely used in this ITSO project and we recommend that you use it when planning the Virtualization Engine installation because:

- It provides you with the latest technical information, such as operating systems support, hardware requirements, documentation, performance recommendations, and so forth.
- It allows you to build an image of your topology and compare different scenarios that you can discuss among professionals before making the final implementation decision.
- It provides you the logical steps of the installation, so that you avoid missing a task during the process.

© Copyright IBM Corp. 2005. All rights reserved.

4.1 Step 1: Accessing the planning advisor

We recommend you start planning and ordering the Virtualization Engine from the planning advisor available, which is available at:

http://publib.boulder.ibm.com/infocenter/eserver/v1r1/en_US/index.htm?info/veinfo/planner.htm

Figure 4-1 shows the first panel.

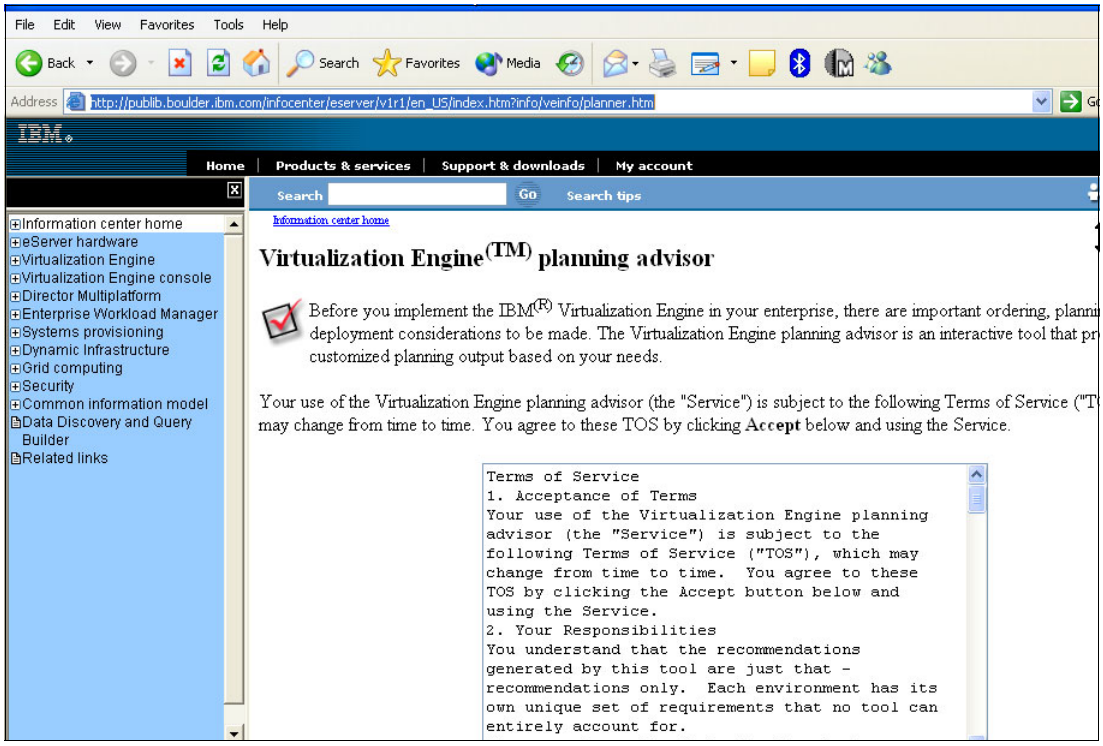
File Edit View Favorites Tools Help
Back Search Favorites Media
Address http://publib.boulder.ibm.com/infocenter/eserver/v1r1/en_US/index.htm?info/veinfo/planner.htm
IBM
Home | Products & services | Support & downloads | My account
Search Go Search tips
Information center home
Information center home
eServer hardware
Virtualization Engine
Virtualization Engine console
Director Multiplatform
Enterprise Workload Manager
Systems provisioning
Dynamic Infrastructure
Grid computing
Security
Common information model
Data Discovery and Query Builder
Related links
Virtualization Engine(TM) planning advisor
Before you implement the IBM(R) Virtualization Engine in your enterprise, there are important ordering, planni
deployment considerations to be made. The Virtualization Engine planning advisor is an interactive tool that pr
customized planning output based on your needs.
Your use of the Virtualization Engine planning advisor (the "Service") is subject to the following Terms of Service ("T
may change from time to time. You agree to these TOS by clicking Accept below and using the Service.
Terms of Service
1. Acceptance of Terms
Your use of the Virtualization Engine planning advisor (the "Service") is subject to the following Terms of Service ("TOS"), which may change from time to time. You agree to these TOS by clicking the Accept button below and using the Service.
2. Your Responsibilities
You understand that the recommendations generated by this tool are just that - recommendations only. Each environment has its own unique set of requirements that no tool can entirely account for.

Figure 4-1 Virtualization Engine planning advisor

The planning advisor is also accessible from the Information Center at:

http://publib.boulder.ibm.com/infocenter/eserver/v1r1/en_US/index.htm?info/icmain.htm

When you have accessed the planning advisor, then choose **Virtualization Engine** → **Systems services** → **Planning**, as shown in Figure 4-2.

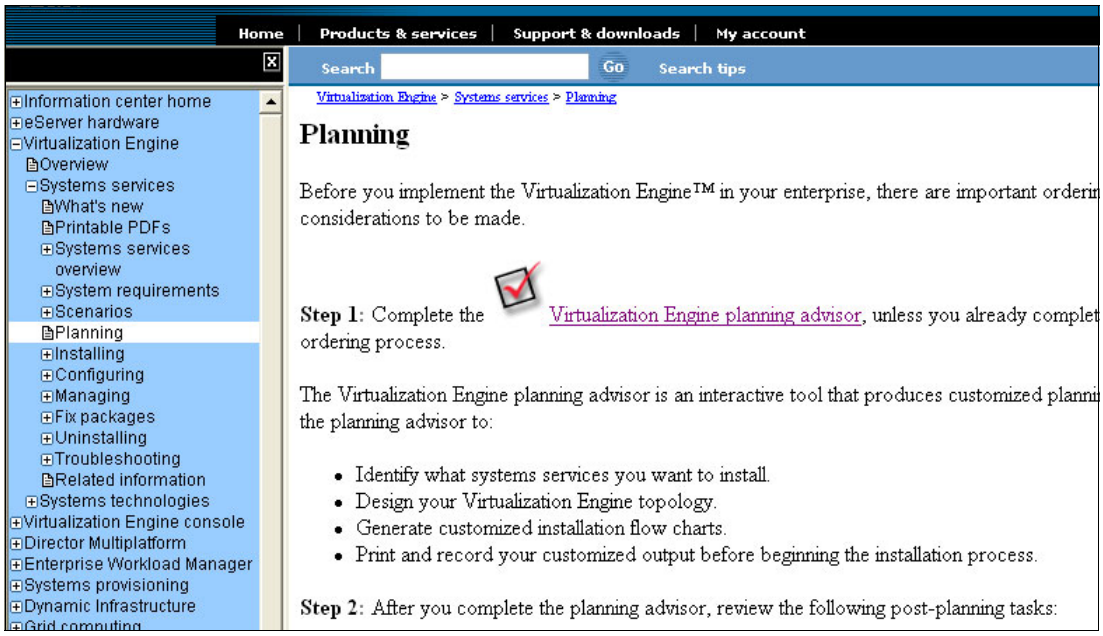

Figure 4-2 Accessing the planning advisor from the Information Center

You can also access the planning advisor from the Virtualization Engine Web site under the planning advisor link, as shown in Figure 4-3:

`http://www-1.ibm.com/servers/eserver/about/virtualization/getting_started.html`

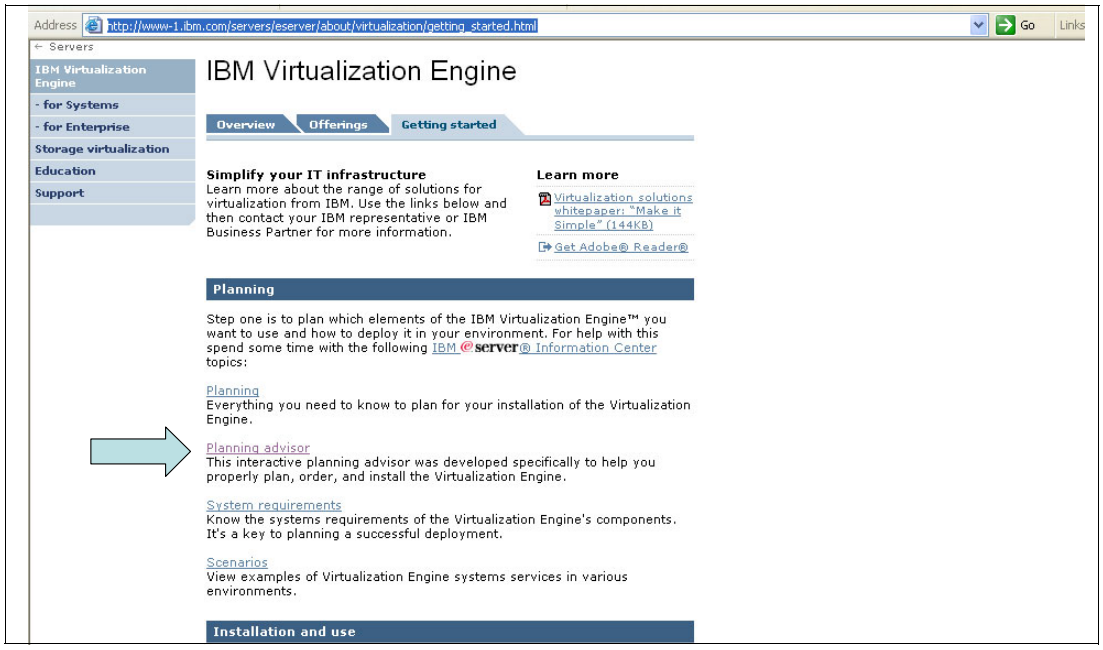
Address
http://www-1.ibm.com/servers/eserver/about/virtualization/getting_started.html
Go
Links
Servers
IBM Virtualization Engine
- for Systems
- for Enterprise
Storage virtualization
Education
Support
IBM Virtualization Engine
Overview
Offerings
Getting started
Simplify your IT infrastructure
Learn more about the range of solutions for virtualization from IBM. Use the links below and then contact your IBM representative or IBM Business Partner for more information.
Learn more
Virtualization solutions whitepaper: "Make it Simple" (144KB)
Get Adobe® Reader®
Planning
Step one is to plan which elements of the IBM Virtualization Engine™ you want to use and how to deploy it in your environment. For help with this spend some time with the following IBM @server® Information Center topics:
Planning
Everything you need to know to plan for your installation of the Virtualization Engine.
Planning advisor
This interactive planning advisor was developed specifically to help you properly plan, order, and install the Virtualization Engine.
System requirements
Know the systems requirements of the Virtualization Engine's components. It's a key to planning a successful deployment.
Scenarios
View examples of Virtualization Engine systems services in various environments.
Installation and use

Figure 4-3 Accessing the planning advisor from the virtualisation engine Web site

After you accept the agreement, if this is the first time you have used the planing advisor or if you did not have saved your data from a previous interview, you see the Welcome panel, as shown in Figure 4-4.

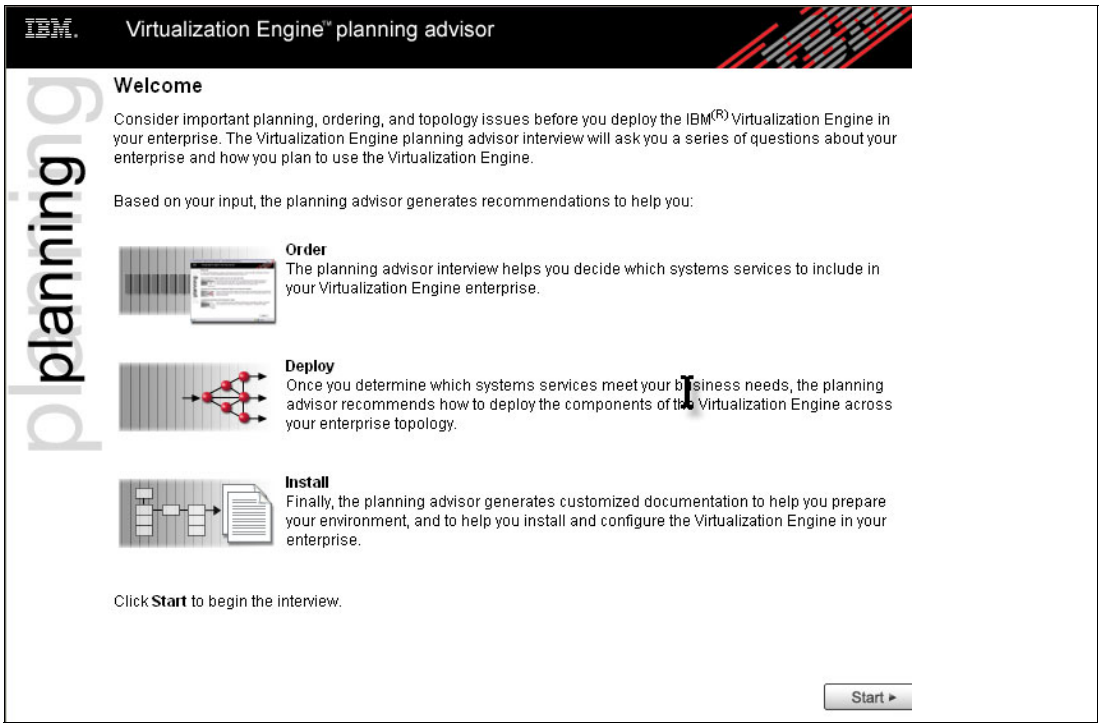
IBM.
Virtualization Engine™ planning advisor
planning
Welcome
Consider important planning, ordering, and topology issues before you deploy the IBM(R) Virtualization Engine in your enterprise. The Virtualization Engine planning advisor interview will ask you a series of questions about your enterprise and how you plan to use the Virtualization Engine.
Based on your input, the planning advisor generates recommendations to help you:
Order
The planning advisor interview helps you decide which systems services to include in your Virtualization Engine enterprise.
Deploy
Once you determine which systems services meet your business needs, the planning advisor recommends how to deploy the components of the Virtualization Engine across your enterprise topology.
Install
Finally, the planning advisor generates customized documentation to help you prepare your environment, and to help you install and configure the Virtualization Engine in your enterprise.
Click Start to begin the interview.
Start ▸

Figure 4-4 Welcome panel from the planning advisor

You can save your interview data after you generate the recommendations by doing one of the following:

- Clicking **Save recommendations** (shown in Figure 4-8 on page 72).
- Clicking **Print this page** at the top of each information tab.
- Closing the planning advisor. When you open it again, you are asked if you would like to use the recommendations that you generated previously.

The next time you enter the planning advisor, you are offered either to use this file or to restart a new configuration.

4.2 Step 2: Completing the interview

The Virtualization Engine planning advisor interview asks you a series of questions about your enterprise and how you plan to use the Virtualization Engine. During the interview, the planning advisor asks which Virtualization Engine services you want to install and on which operating systems, as shown in Figure 4-5. The planning advisor is configured and updated to give you the latest information regarding the possible topology and the requirements. For example, the planning advisor includes the latest available agents, the latest available fix packs, and the latest recommendations from the latest experiences, according to your input.

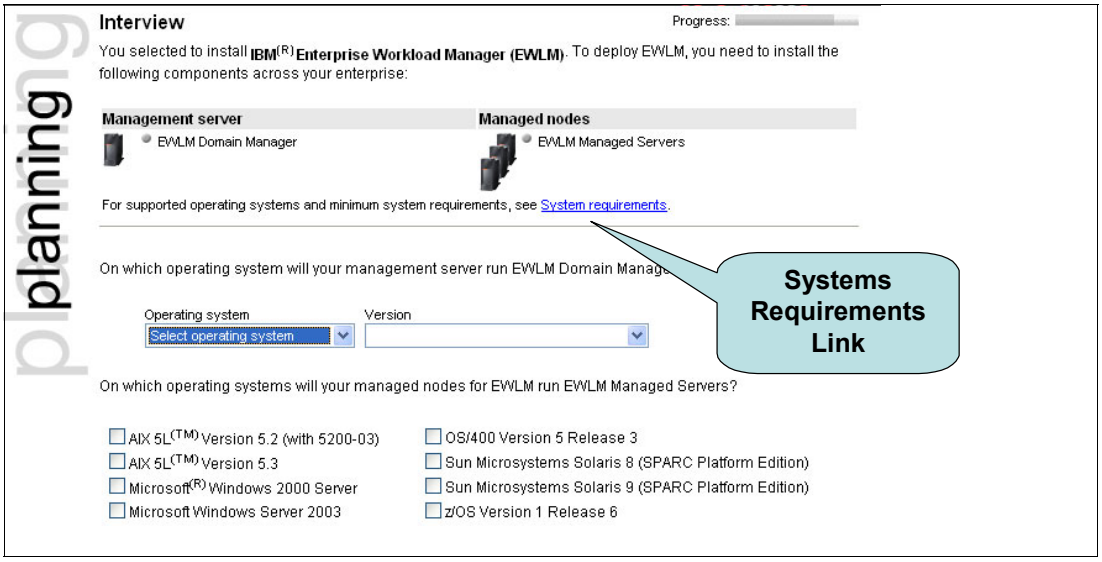
planning
Interview
Progress:
You selected to install IBM(R) Enterprise Workload Manager (EWLM). To deploy EWLM, you need to install the following components across your enterprise:
Management server
EWLM Domain Manager
Managed nodes
EWLM Managed Servers
For supported operating systems and minimum system requirements, see System requirements.
On which operating system will your management server run EWLM Domain Manage
Operating system
Select operating system
Version
Systems Requirements Link
On which operating systems will your managed nodes for EWLM run EWLM Managed Servers?
AIX 5L(TM) Version 5.2 (with 5200-03)
AIX 5L(TM) Version 5.3
Microsoft(R) Windows 2000 Server
Microsoft Windows Server 2003
OS/400 Version 5 Release 3
Sun Microsystems Solaris 8 (SPARC Platform Edition)
Sun Microsystems Solaris 9 (SPARC Platform Edition)
z/OS Version 1 Release 6

Figure 4-5 Planning advisor interview

For the interview, you need to understand the topology definitions of the Virtualization Engine. You will be asked to decide which managed servers you will use. The planning advisor also includes a useful link for systems requirements so that you can have an idea about the minimum size of the server that you need to provide. If you select this link, the planning advisor displays the general guidelines for EWLM minimum system requirements, as shown in Figure 4-6.

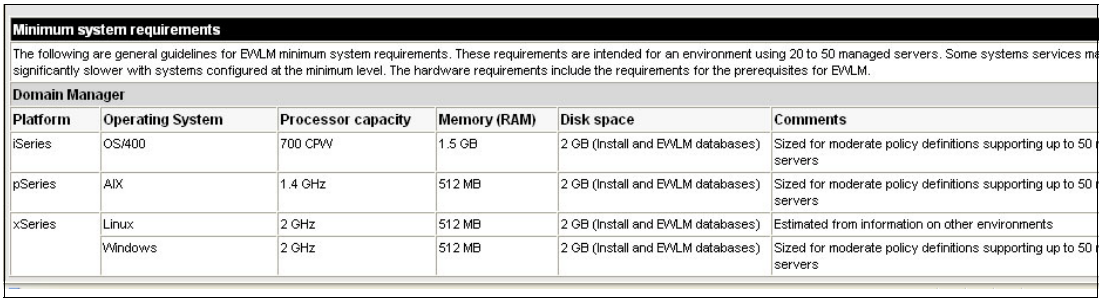

Minimum system requirements

The following are general guidelines for EWLM minimum system requirements. These requirements are intended for an environment using 20 to 50 managed servers. Some systems services m
significantly slower with systems configured at the minimum level. The hardware requirements include the requirements for the prerequisites for EWLM.

Domain Manager

Platform	Operating System	Processor capacity	Memory (RAM)	Disk space	Comments
iSeries	OS/400	700 CPW	1.5 GB	2 GB (Install and EWLM databases)	Sized for moderate policy definitions supporting up to 50 servers
pSeries	AIX	1.4 GHz	512 MB	2 GB (Install and EWLM databases)	Sized for moderate policy definitions supporting up to 50 servers
xSeries	Linux	2 GHz	512 MB	2 GB (Install and EWLM databases)	Estimated from information on other environments
	Windows	2 GHz	512 MB	2 GB (Install and EWLM databases)	Sized for moderate policy definitions supporting up to 50 servers

Figure 4-6 Systems requirements

The last step of the interview is a summary of your configuration, from where you generate recommendations, as shown in Figure 4-7.

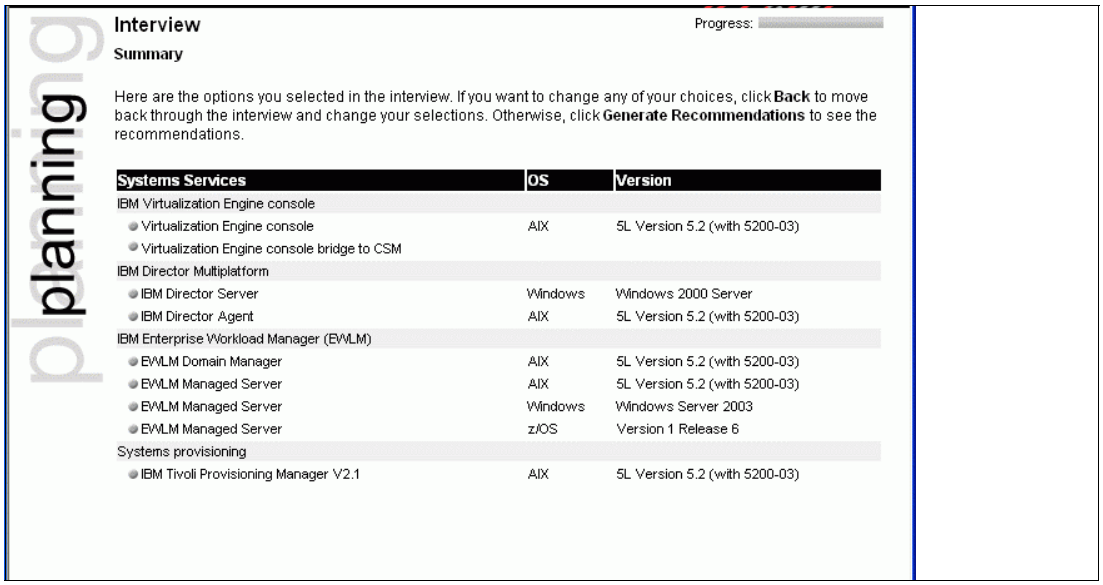

Figure 4-7 Interview summary

When you generate the recommendations, you leave the interview phase to enter into the recommendations phase, as shown in Figure 4-8.

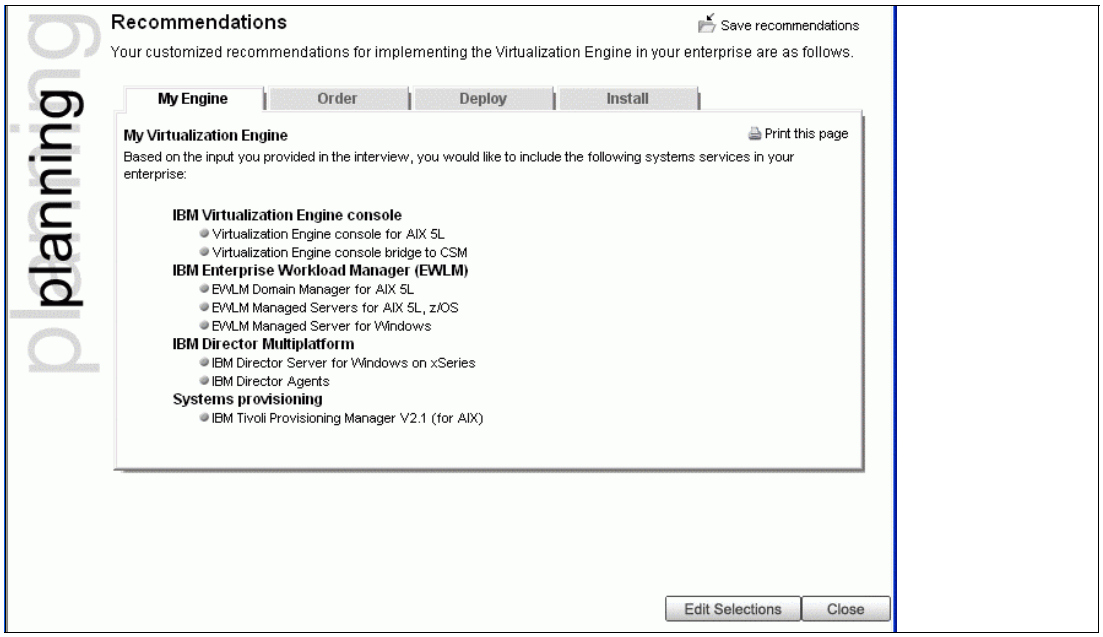
planning
Recommendations
Save recommendations
Your customized recommendations for implementing the Virtualization Engine in your enterprise are as follows.
My Engine
Order
Deploy
Install
My Virtualization Engine
Print this page
Based on the input you provided in the interview, you would like to include the following systems services in your enterprise:
IBM Virtualization Engine console
Virtualization Engine console for AIX 5L
Virtualization Engine console bridge to CSM
IBM Enterprise Workload Manager (EWLM)
EWLM Domain Manager for AIX 5L
EWLM Managed Servers for AIX 5L, z/OS
EWLM Managed Server for Windows
IBM Director Multiplatform
IBM Director Server for Windows on xSeries
IBM Director Agents
Systems provisioning
IBM Tivoli Provisioning Manager V2.1 (for AIX)
Edit Selections
Close

Figure 4-8 Recommendations

Click **Edit Selections** to return to the interview panel. The Recommendations panel shows the recommendations profile.

4.3 Step 3: Ordering the code

The Order tab on the Recommendations panel provides information how to get the code that you need.

4.3.1 Passport Advantage Online

The Virtualization Engine is only available through Passport Advantage® Online at:

http://www.ibm.com/software/passportadvantage

Passport Advantage Online enables you to manage your account by downloading software, purchasing new software licenses, renewing Software Maintenance coverage or Fixed Term Licenses, generating reports, updating account information, viewing or printing Proofs of Entitlement, and much more. To access Passport Advantage Online, customers and IBM Business Partners are required to have an IBM registration ID and password, which can be created by accessing Passport Advantage Online at:

http://www.ibm.com/software/passportadvantage

Select the For customers or For business partners tab and then the appropriate sign in link. To create a new IBM registration ID and password, select the register link and complete the necessary information.

This suite can be ordered in two packages:

- The 5724-i71 code is an a la carte offering, from which you can pick up any of the Virtualization Engine suite components.
- The 5724-i72 code allows you to purchase the full set of components depending on the platform you want to install these components.

IBM Virtualization Engine Management Servers a la carte offering: 5724-i71

You can order elements from the Virtualization Engine suite as separate, optional components as described in the following lists:

- Virtualization Engine Enterprise Workload Manager for managing AIX 5L and i5/OS which includes:
  - IBM Enterprise Workload Manager agents for AIX 5L and i5/OS
  - Virtualization Engine V1.1 installer, console, and common runtime
  - IBM Integrated Solutions Console 5.0.2
  - WebSphere Application Server V5.1
  - IBM Tivoli Directory Server V5.1
- Virtualization Engine Enterprise Workload Manager for managing Microsoft Windows and Sun Solaris which includes:
  - Enterprise Workload Manager agents for Microsoft Windows and Sun Solaris 8 and 9 (SPARC Platform Edition)
  - Virtualization Engine V1.1 installer, console, and common runtime
  - Integrated Solutions Console 5.0.2
  - WebSphere Application Server V5.1
  - IBM Tivoli Directory Server V5.1

- Virtualization Engine IBM Director which includes:
  - Director Multiplatform
  - Virtualization Engine V1.1 installer, console, and common runtime
  - Integrated Solutions Console 5.0.2
  - WebSphere Application Server V5.1
  - IBM Tivoli Directory Server V5.1
  - DB2 UDB Workgroup Unlimited Edition 8.1
  - Grid Toolbox V3 for Multiplatforms

IBM Virtualization Engine Services Multiplatform integrated offering: 5724-i72

The Virtualization Engine suite code 5724-i72 is available in two options, depending upon the choice of IBM Enterprise Workload Manager-managed servers and agents. These IBM Enterprise Workload Manager managed servers and agents come in two separate packages:

- Virtualization Engine systems services for managing AIX 5L and i5/OS which includes:
  - Enterprise Workload Manager domain manager
  - Enterprise Workload Manager managed servers and agents for AIX 5L and i5/OS
  - IBM Director Multiplatform management server and agents
  - IBM Tivoli Provisioning Manager
  - Grid Toolbox V3 for Multiplatforms
  - Virtualization Engine installer, Virtualization Engine console, and common runtime infrastructure
- Virtualization Engine systems services for managing Microsoft Windows and Sun Solaris which includes:
  - Enterprise Workload Manager domain manager
  - Enterprise Workload Manager managed servers and agents for Microsoft Windows and Sun Solaris 8 and 9 (SPARC Platform Edition)
  - IBM Director Multiplatform management server and agents
  - Tivoli Provisioning Manager
  - Grid Toolbox V3 for Multiplatforms
  - Virtualization Engine installer, Virtualization Engine console, and common runtime infrastructure

4.3.2 Important notes

The following lists important notes for your consideration:

- The IBM Director component, can be ordered separately from the Virtualization Engine suite. However, when ordered through another channel, IBM Director does not include IBM Software Maintenance, which is only available through the Virtualization Engine.
- You can get Tivoli Provisioning Manager either from the IBM Virtualization Engine Services Multiplatform (5724-i72) or from the IBM software Tivoli catalog under the product code 5724-i15. The code delivered is the same.
- iSeries has a specific offering, 5733-VE1, which allows you to get some components of the Virtualization Engine when you order an IBM selected iSeries Editions (520, 550, or 570). The IBM Virtualization Engine for iSeries (5733-VE1) includes the IBM Director on all V5R3 iSeries operating system orders and IBM Enterprise Workload Manager on selected iSeries editions.
- IBM Virtualization Engine Enterprise Workload Manager for z/OS V1.1.0 (EWLM for z/OS, for managed nodes) is available. This initial release provides the capability for z/OS to be a managed server for monitoring z/OS applications within EWLM. You can order it using

the product ID 5665-M76. It is a $0-priced product and it is distributed in one package and contains the basic machine readable material, the program directory, and the license information.

- Initial offerings of IBM Grid Toolbox V3 for Multiplatforms are available on AIX 5L for pSeries, and Linux on xSeries.
  - 5765-G22 IBM Grid Toolbox V3 for AIX is offered at no charge.
  - IBM Grid Toolbox V3 for Linux on xSeries is offered as no charge. 52P1214 is the part number to use if you want a product media package. 52P1263 is the part number to use for electronic download.

Note: If you are an i5/OS customer, IBM Director is shipped with every i5/OS order.

Note: When ordering the IBM Director Multiplatform product, you also receive the Virtualization Engine console and the Virtualization Engine common runtime.

4.3.3 Summary

Table 4-1 and Table 4-2 summarize all the product codes that deliver parts of the Virtualization Engine and which product to use to get Virtualization Engine or one component.

Table 4-1 By product number

This code	Allows you to get the following
5724-i71	Any of the Virtualization Engine suite components
5724-i72	All the Virtualization Engine suite components together
5733-VE1	► EWLM with the iSeries ► IBM Director Multiplatform with the iSeries
5685-M76	EWLM Managed Server on z/OS
5724-i15	Tivoli Provisioning Manager through Tivoli
5765-G22	IBM Grid Toolbox

Table 4-2 By component

If you want:	Use the following
EWLM	5724-i71 or 5724-i72
EWLM on iSeries	5724-i71, 5724-i72, or 5733-VE1
EWLM Managed Server on the zSeries	5685-M76
Tivoli Provisioning Manager	5724-i15, 5724-i71, or 5724-i72
Virtualization Engine console	5724-i71 or 5724-i72
IBM Director Multiplatform	5724-i71 or 5724-i72

4.4 Step 4: Obtaining the code

The planning advisor guides you through how to get the latest available code. Different channels exist, depending whether you need the basic code or the fix packs.

4.4.1 Basic code

There are two methods to get the Virtualization Engine code from Passport Advantage:

- CD: You receive the installation files on physical media (CDs). Depending on the components and the options that you choose, you might receive multiple CDs.
- Download: You can download the necessary installation files from the Passport Advantage Web site directly to a server.

Note: Tivoli Provisioning Manager for Solaris is available as a download only. You can still receive the remainder of your Virtualization Engine software on CDs, but you have to download Tivoli Provisioning Manager for Solaris.

For this project, we use files that we downloaded and we put these files on an FTP server so that different teams were able to share and use the code.

4.4.2 Fix packs

You also need to look at the latest available fix packs. *Fix packs* are updates to software that fix known problems or add new functions. You apply fix packs to keep hardware, software, and operating systems current. There is an installable fix pack for each Virtualization Engine systems service on each operating system. For example, there is a fix pack for managed server on AIX and another fix pack for EWLM Managed Server on i5/OS. A fix pack includes code for all fix packs that are available for the release. Thus, you do not need to install multiple fix packs. You can just install the latest fix pack.

Fix packs are named in the form of *v.r.f*:

- The *v* and *r* correspond to the Virtualization Engine version number. For example, all fix packs for Virtualization Engine Version 1.1 begin with *1.1*.
- The *f* indicates the fix pack number. Version numbers are ordered but are not necessarily sequential.

You can download fix packs at:

`http://techsupport.services.ibm.com/server/VirtualizationEngine`

As shown on Figure 4-9, you can search the available fix packs for a particular operating system, for a particular Virtualization Engine systems service, or by group.

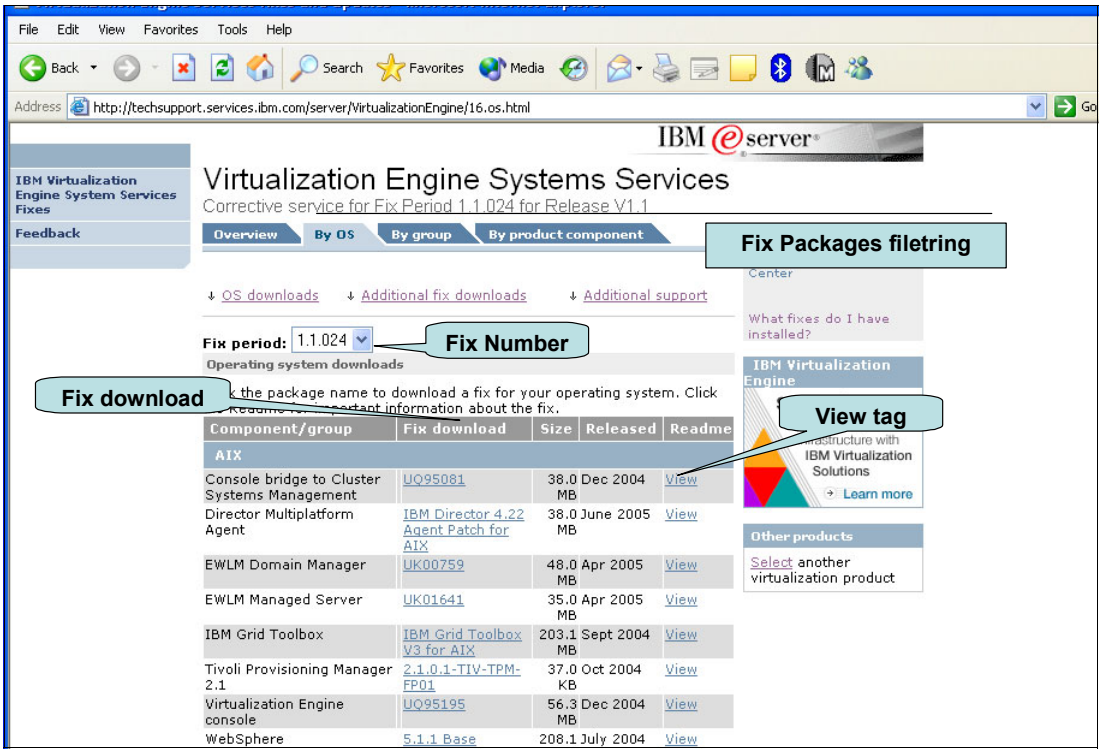

Figure 4-9 IBM Virtualization Engine systems services fix pack Web site

The *fix download number* is an IBM internal number that is unique to the problem. From this Web site, you can click **View** to see a description of the corrections or the functions that are included in the fix pack.

Fix packs are installed using an installation wizard, which is provided as an executable file within the archived fix pack that you download. The fix pack installation process is similar to the installation process for the Virtualization Engine itself. You can uninstall fix packs.

Tip: The recommendation is to install all the fix packs that are available for your environment.

4.5 Step 5: Understanding your logical topology

Through the Deploy tab, as shown in Figure 4-10, the planning advisor proposes a logical topology of your specific Virtualization Engine environment, from which you can get more detailed information by server or by component.

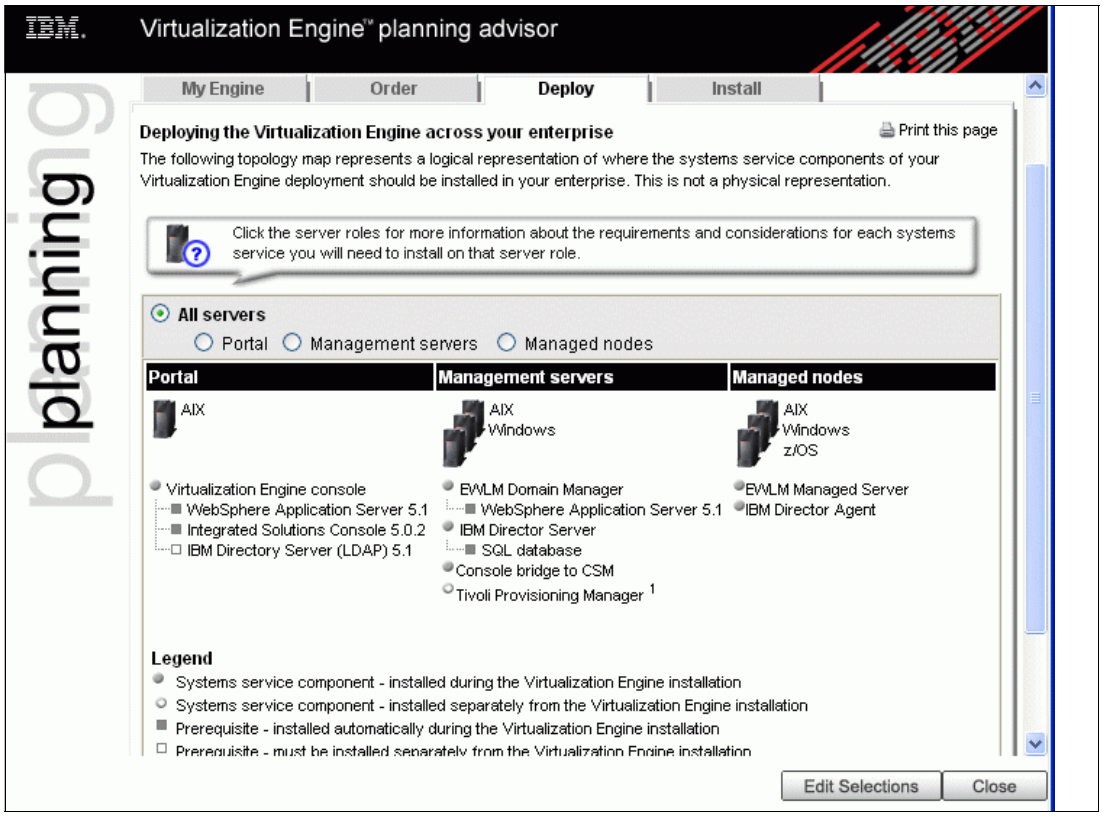
IBM.
Virtualization Engine™ planning advisor
planning
My Engine
Order
Deploy
Install
Deploying the Virtualization Engine across your enterprise
Print this page
The following topology map represents a logical representation of where the systems service components of your Virtualization Engine deployment should be installed in your enterprise. This is not a physical representation.
Click the server roles for more information about the requirements and considerations for each systems service you will need to install on that server role.
All servers
Portal
Management servers
Managed nodes
Portal
Management servers
Managed nodes
AIX
AIX
Windows
AIX
Windows
z/OS
Virtualization Engine console
WebSphere Application Server 5.1
Integrated Solutions Console 5.0.2
IBM Directory Server (LDAP) 5.1
EWLM Domain Manager
WebSphere Application Server 5.1
IBM Director Server
SQL database
Console bridge to CSM
Tivoli Provisioning Manager 1
EWLM Managed Server
IBM Director Agent
Legend
Systems service component - installed during the Virtualization Engine installation
Systems service component - installed separately from the Virtualization Engine installation
Prerequisite - installed automatically during the Virtualization Engine installation
Prerequisite - must be installed separately from the Virtualization Engine installation
Edit Selections
Close

Figure 4-10 Logical Topology from the planning advisor

This view provides very important information to understand better which servers need to host the Virtualization Engine components, how these servers are installed (individually or not), and the prerequisites. This view is a *logical* view and not a *physical* view.

When you have a good understanding of the Virtualization Engine topology, the components, and their relationship, you can draw such a topology by yourself from the beginning. However, we recommend that you use the planning advisor and check your topology with the topology that the planning advisor proposes. The planning advisor is updated regularly. It might reflect changes of which you are not aware or might propose other ways to build your logical topology.

4.6 Step 6: Matching the logical with the physical topology

At this point, you are ready to match the hardware topology that you have chosen to host the Virtualization Engine with the deployment topology proposed from the previous step. To do so, you must consider your current enterprise environment. The following questions will arise, and you will need the answers before you begin the installation process:

- Identify the security environment: What firewalls do you need to go through? Which authorizations do you need? Which users need authorization?
- Identify the network environment: Are you using a domain name server? Virtual LANs? Subnets? Are IP addresses assigned dynamically?
- Identify the storage infrastructure that you can use: How are servers connected to the storage infrastructure?
- Identify if you already have a directory server installed and configured in your topology and if it can be used for Virtualization Engine.

  Attention: the Directory server for the Virtualization Engine must be the Tivoli Directory Server product at a specific level that depends on the server.

- Identify if a database from your environment can be used for the topology.
- Identify possible product co-existence and runtime considerations.

Other questions might arise, depending on your objectives, your infrastructure, and your skills.

4.7 Step 7: Learning about the installation process

The Install tab, as shown in Figure 4-11, proposes a series of installation tasks. These tasks are structured by Virtualization Engine component. Clicking an installation task shows a detailed list of subtasks, such as prerequisites, installation, post-installation, configuration, and post-configuration tasks.

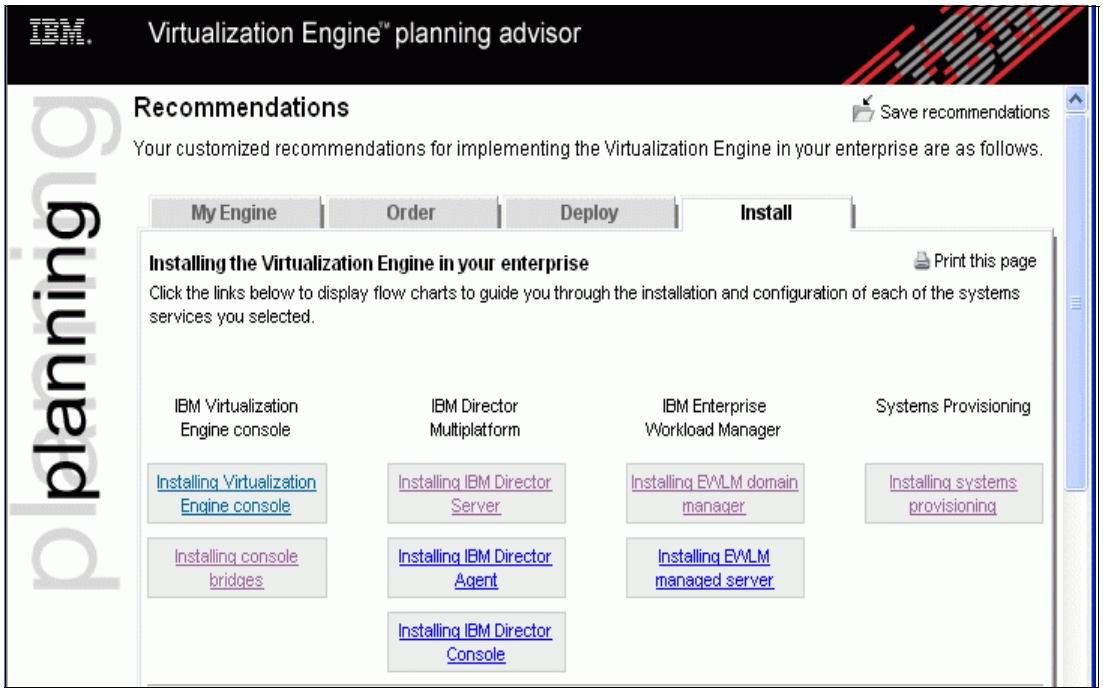
IBM.
Virtualization Engine™ planning advisor
planning
Recommendations
Save recommendations
Your customized recommendations for implementing the Virtualization Engine in your enterprise are as follows.
My Engine
Order
Deploy
Install
Installing the Virtualization Engine in your enterprise
Print this page
Click the links below to display flow charts to guide you through the installation and configuration of each of the systems services you selected.
IBM Virtualization Engine console
IBM Director Multiplatform
IBM Enterprise Workload Manager
Systems Provisioning
Installing Virtualization Engine console
Installing IBM Director Server
Installing EWLM domain manager
Installing systems provisioning
Installing console bridges
Installing IBM Director Agent
Installing EWLM managed server
Installing IBM Director Console

Figure 4-11 Installation tasks by component

The planning advisor identifies the difficulty level of each of these subtasks with a specific icon. However, the difficulty depends on your skill set, as shown in Figure 4-12 on page 81.

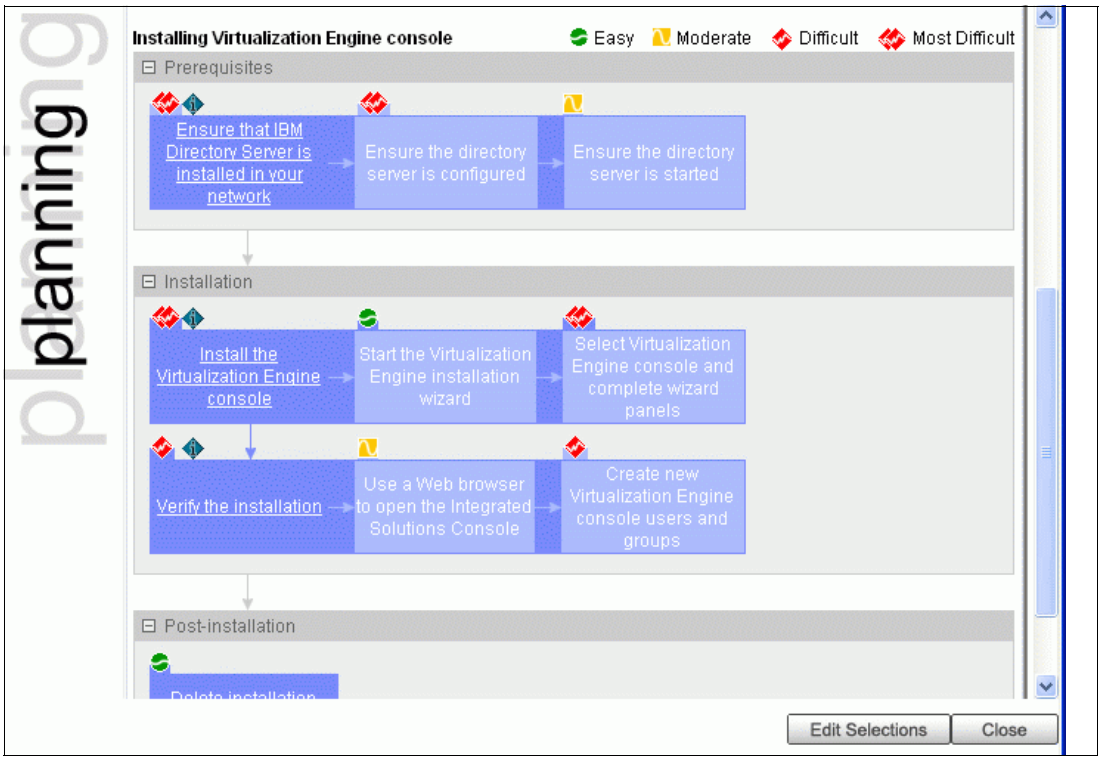

Figure 4-12 Installation subtasks

Clicking a subtask forwards you to the Information Center for details about how to complete that subtask.

The wizard installs EWLM, IBM Director, the Virtualization Engine console, and the common runtime components to support these products. Grid Toolbox V3 for Multiplatforms and systems provisioning that is delivered by Tivoli Provisioning Manager are installed using separate wizards.

4.8 Summary

The planning advisor is a key tool when planning to install the Virtualization Engine. In addition, we recommend that you use the worksheets that are provided through the Information Center. You should also consult the available documentation before you begin the installation process.

Part 2

Installing the Virtualization Engine

The chapters in this part of the redbook describe the installation process for the Virtualization Engine components. Because you can install the Virtualization Engine on different operating systems, different scenarios are available . However, these scenarios do not cover all the current possibilities of the Virtualization Engine. You might choose other scenarios. We choose the scenarios that we describe here according to customer interest that we have learned from experience. Another guidance was to develop scenarios using a common set of preferred management servers. We recommend that you use one management server because the installation is simplified if the set of management servers uses the same type of operating system.

The scenarios that we choose are:

- Chapter 5, “Installing IBM Director, the Virtualization Engine console, and TPM on a BladeCenter with xLinux” on page 85
- Chapter 6, “Installing the Virtualization Engine console, IBM Director, TPM, and EWLM on pSeries with AIX” on page 117
- Chapter 7, “Installing the Virtualization Engine console, IBM Director, TPM, and EWLM on Blade with Windows” on page 161
- Chapter 9, “Installing Virtual Machine Manager and VMware on managed servers” on page 225

We did not cover in this project a scenario with an iSeries as a management server. However Appendix 8, “Installing and configuring the Virtualization Engine for IT management in an OS/400 environment” on page 201 describes an example of the installation process in a iSeries environment. You can find this scenario in the Information Center.

© Copyright IBM Corp. 2005. All rights reserved.

The installation processes might be redundant across the platforms. However, different teams build these scenarios in different ways, depending on skill, knowledge, and work habits. You might find some interest in reading the different test cases.

5

Installing IBM Director, the Virtualization Engine console, and TPM on a BladeCenter with xLinux

This chapter describes the various installation and configuration steps that we used for managing the infrastructure from an xSeries for Linux environment for the following components:

- IBM Director
- Virtualization Engine console
- IBM Tivoli Provisioning Manager

This chapter discusses the following topics:

- Before the installation
- The installation process
- Fix pack installation considerations
- Post-installation step: configuring the data center model

© Copyright IBM Corp. 2005. All rights reserved.

5.1 Before the installation

Before the actual installation, we describe the topology, the installation flow, and the few pre-installation steps that you have to consider.

5.1.1 The topology

In this topology, we use two xSeries Linux machines as our management servers to manage the workload that runs across multiple platforms. In this topology:

- The first xSeries Linux machine is used to configure the Virtualization Engine console and the IBM Director that the Virtualization Engine console uses to obtain its information for managing the environment, otherwise known as its *management source*.
- The second xSeries Linux machine is used for Tivoli Provisioning Manager for provisioning servers into and out of our IT environment, based on our business needs.

We used two Blades of an HS20-8832 server, each Blade had four processors. The two Linux for xSeries images were SuSE Linux version 2.4.21-278-smp.

Figure 5-1 shows the management topology that we used with two management servers.

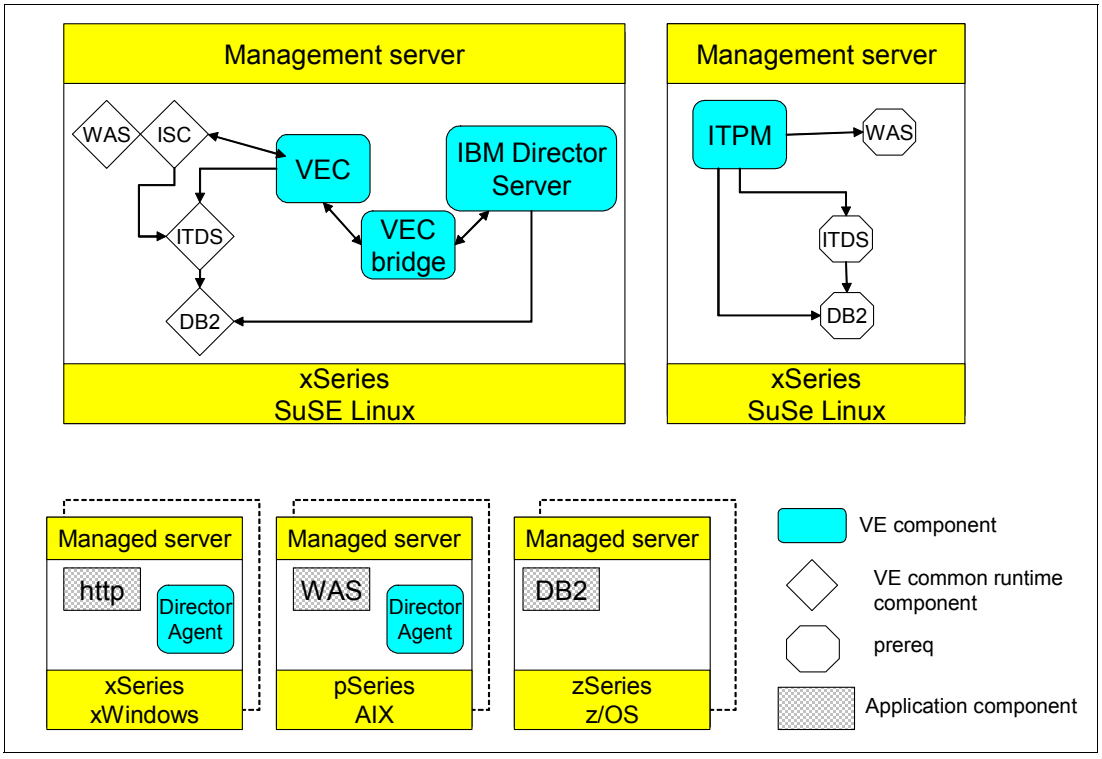

Figure 5-1 Topology for the xLinux environment

5.1.2 The installation flow

In this configuration, the main steps of the installation are the following:

1. Installation and configuration of the prerequisite components, which include the following:
   a. The Tivoli Directory Server, which is the required Lightweight Directory Access Protocol (LDAP) server for the Integrated Solutions Console (ISC) in which the code runs.
   b. The database that is required by the LDAP server, DB2 in our case.
   c. A WebSphere Application Server, which is needed to run the administrative application which is a Java application and run the ISC.
2. Installation of the Virtualization Engine console, the IBM Director, and the bridges to link the Virtualization Engine console and the IBM Director.
3. Installation of the IBM Director agents on the managed servers.
4. Installation of Tivoli Provisioning Manager on the management server which needs a WebSphere Application Server (to run the administrative application), a database server (DB2 Universal Database Workgroup Server Unlimited Edition V8.1 with Fix Pack 3) and an LDAP server (Tivoli Directory Server in our case), which, in turn, needs a database (DB2 in our case).

The following sections describe these steps in more detail.

5.1.3 Pre-installation recommendations

After careful planning of which components you will use, you can obtain and use the code by following these steps:

1. It is recommended that you download the Virtualization Engine files from Passport Advantage to a central location and make the files available to the other systems in your environment. If possible, these tar (or zipped) files should be put on an NFS server so that they can be accessed by each of the management and managed servers.
2. It is recommended that you have the latest Virtualization Engine fix packs available (see “Fix pack installation considerations” on page 109). Before installing the components, you should check what fix packs are available, what problems have been resolve, and what workarounds exist.

   For example, the installation of DB2 using the graphical installer fails due to a port conflict when installing Tivoli Provisioning Manager on xLinux SuSE 8. The graphical installation process selects port number 50000 by default. However, this port is already assigned by the system. There is a workaround documented for this problem in the Tivoli Intelligent Orchestrator/Tivoli Provision Manager Fix Pack 01.
3. It is recommended that you use the Information Center as the primary source for installation information for Virtualization Engine because it contains the most up-to-date installation information that is available. It also contains PDF files for each of the components, which can be viewed online or printed.

To assist with the installation of the components into our environment, we use the planning advisor in the following ways:

- To help determine which systems services we need to include.
- To deploy these systems services into our environment.
- To provide the installation information, which includes the following:
  - Software prerequisites (such as Tivoli Directory Server, WebSphere Application Server, and DB2).
  - Installation of the Virtualization Engine systems services (such as the Virtualization Engine console).
  - Post-installation processes.
  - Configuration and optional post-configuration steps.
  - Installation verification.

This information is obtained by clicking the links within the planning advisor, which direct you to the Information Center.

Tips:

- We found the planning advisor hyperlinks were *very helpful* because they brought us quickly to the correct information in the Information Center.
- Additionally, printing the documentation for each of the components that we were installing was also very useful because it allowed us to complete the worksheets and to check off each step as we completed it.
- We also used these worksheets as a reference when installing the next component. For example, the Virtualization Engine console and the IBM Director use some of the same definitions.

Let us talk now about how we use the planning advisor to help us with our installation for the Virtualization Engine console, IBM Director, and Tivoli Provisioning Manager.

5.1.4 Using the Virtualization Engine planning advisor

While using the planning advisor, we specify that we would like to include IBM Director, systems provisioning, and the Virtualization Engine console system services using the following:

1. We want a single interface for system administrators to monitor and manage multiple platforms and servers. We specify that we want to install the Virtualization Engine console. The Virtualization Engine console provides a Web console that connects to system and storage management applications, such as the IBM Director. Because we want to maximize our xSeries for Linux skills, we specify that we want to install the Virtualization Engine console on xSeries for Linux using SuSE Linux Enterprise Server 8.
2. We need a comprehensive systems-management solution that provides proactive management features across our multiple server types (such as event logs, resource monitors, and thresholds). We specify the IBM Director as our system and storage management application on xSeries for Linux using SuSE Linux Enterprise Server 8, and we specify that the IBM Director Agent runs on the systems where our workload runs (our managed servers). These systems consist of AIX V5.3, SuSE Linux Enterprise Server 8 (IA32 IA64, AMD64), and Windows 2003 EE, SE, and WE.

Note: The Virtualization Engine console connects to the IBM Director and uses it as its management source to manage these resources from a single Web-based location.

3. We need an efficient provisioning mechanism to allocate new resources into our environment when our business goals were not met due to system constraints. Therefore, we specify that we want to use Tivoli Provisioning Manager and want it to run on xSeries for Linux using SuSE Linux Enterprise Server 8.

The following sections describe the most useful tabs in the planning advisor.

Planning advisor Deploy tab

Based on our input, the Virtualization Engine planning advisor provided us with customized recommendations for implementing the Virtualization Engine components. The planning advisor identified the prerequisite software that the components needed, as highlighted by A in Figure 5-2. Highlighted by B in Figure 5-2 is the following information:

- The Portal server lists the Virtualization Engine console along with its prerequisite software products: WebSphere Application Server, Integrated Solutions Console, and IBM Directory Server (LDAP) installed on a Linux machine.
- The Management servers lists the IBM Directory Server with its prerequisite product, the SQL database, along with the Virtualization Engine console bridge and Tivoli Provisioning Manager installed on a Linux machine.

Important: The prerequisite products for Tivoli Provisioning Manager are WebSphere Application Server 5.1.0.3, DB2 V8.1 FP3, and Tivoli Directory Server 5.1

- The managed servers (actually worded as *Managed nodes*) shows the IBM Director Agent installed on AIX, Linux, and Windows.

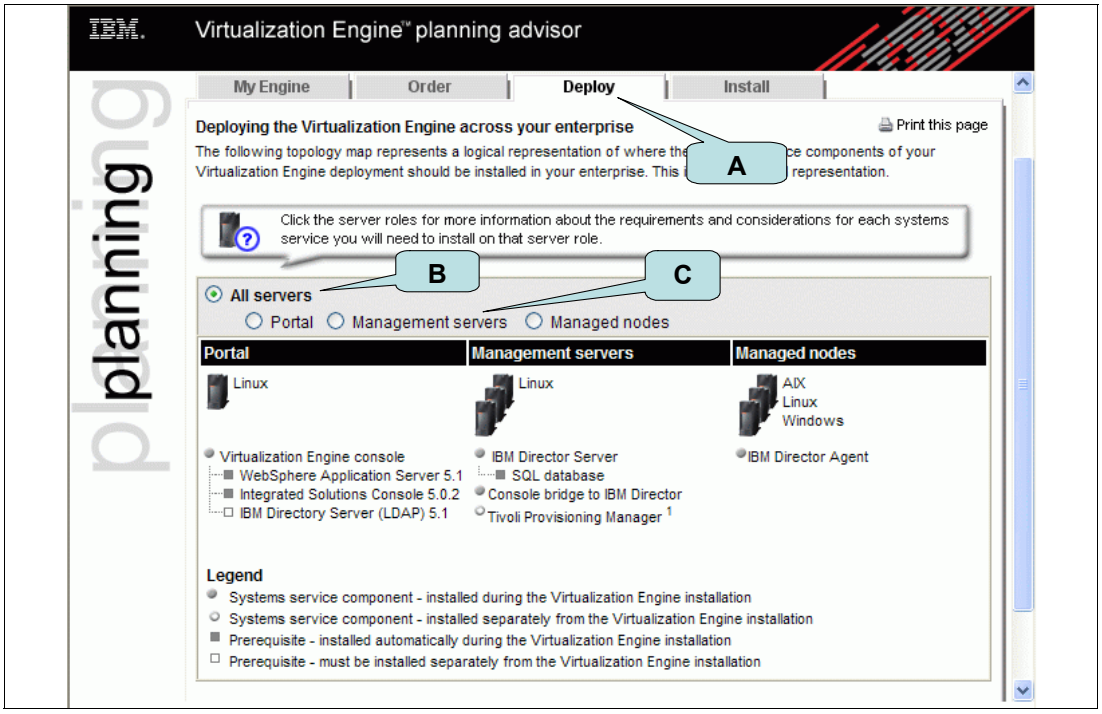

Figure 5-2 The planning advisor Deploy tab (logical mapping)

Additionally, Figure 5-2 on page 89 shows which systems service component and prerequisite gets installed automatically during the Virtualization Engine installation and which must be installed separately as described in the Legend. For example, WebSphere Application Server and ISC are installed automatically during the Virtualization Engine installation, but you must install LDAP and Tivoli Provisioning Manager separately.

This is a *logical* representation of where the systems services should be installed, not a *physical* representation. As stated earlier, our Virtualization Engine console and IBM Director are on one image, while Tivoli Provisioning Manager is on another image.

Note: The information in Figure 5-2 that appears when you select **All servers** (highlighted by B) and when you select one of the server roles (highlighted by C) is unique. The Legend can only be seen when you select **All servers**. However, information such as the minimum system requirements can be obtained by selecting one of the server roles: Portal, Management servers, or managed servers.

Planning advisor Install tab

The last tab in the planning advisor is the Install tab, as highlighted in A in Figure 5-3. This tab lists the components that we chose to install.

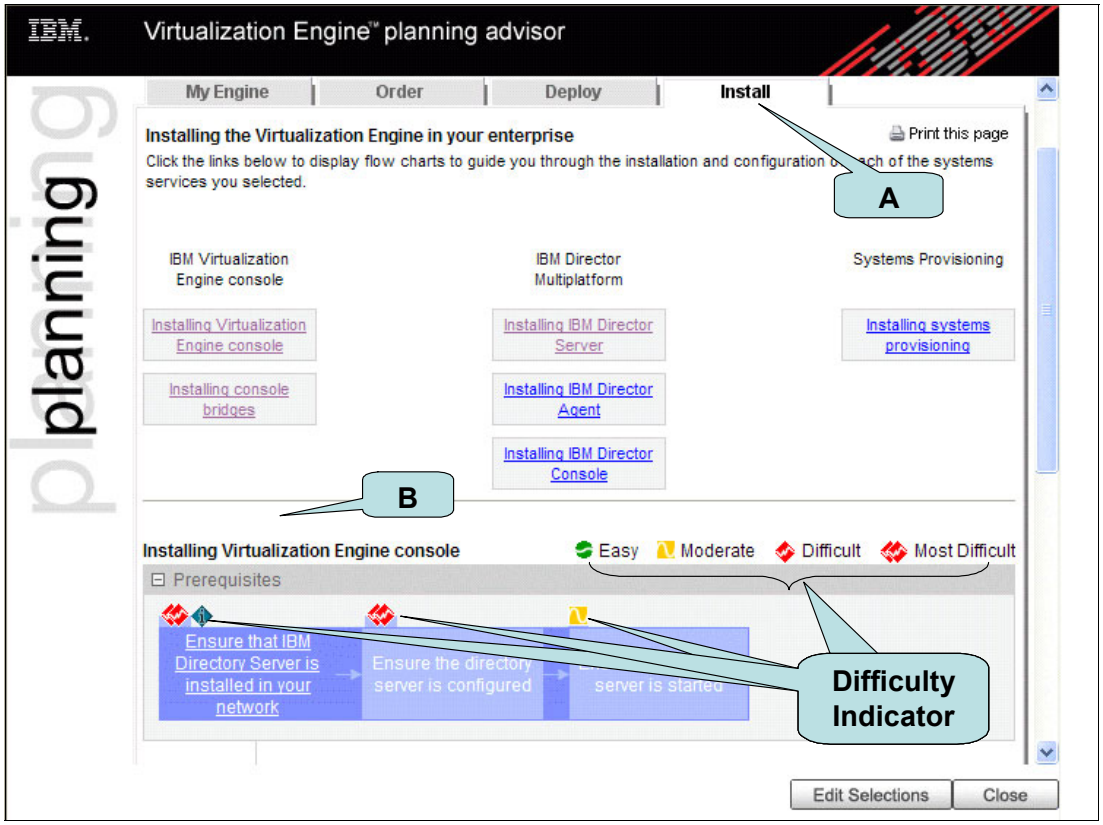

Figure 5-3 The planning advisor Install tab

Using the information from the planning advisor, as shown in Figure 5-3, we installed the system services and their prerequisites in the following order:

1. Virtualization Engine console
2. IBM Director
3. Tivoli Provisioning Manager

When we select any of these components, for example, the Virtualization Engine console, the corresponding flow chart to guide us through the installation and configuration of that component is shown, highlighted by B in Figure 5-3 on page 90. Associated with the steps in the flowchart is a difficulty indicator for performing that function. For example, the first prerequisite for installing the Virtualization Engine console is to ensure that IBM Directory Server is installed in your network.

One of the benefits of using the planning advisor is that we can click the flowchart boxes and link to the Information Center (which appears in another browser window) directly to the section that explains what we need to do. Therefore, when we select **Ensure that IBM Directory Server is installed in your network,** we go to the Information Center instructions for installing the IBM Directory Server. When installed, we can follow the instructions to ensure that the directory server is configured and started. You can find the link at:

```
http://publib.boulder.ibm.com/infocenter/eserver/v1r1/en_US/index.htm?info/esmcinfo/
eicacldapinstall.htm
```

Providing that we have verified previously that we have met all of the requirements for the Virtualization Engine and the Virtualization Engine console as described in “Verifying the requirements” on page 91, we can begin to install the Virtualization Engine console, starting with the prerequisite IBM Directory Server.

Tip: If you exit the planning advisor (for example by closing the window) you do not lose your entries. You can always go back into the Information Center → Virtualization Engine → Systems services → Planning and select **Virtualizing Engine planning advisor**. The planning advisor knows that it has been completed on your computer and after accepting the Terms of Service you are prompted to Load the previous recommendations or create new ones. If you select **Load My Previous Recommendations**, the planning advisor returns all of the values you entered previously, and you can continue with the install.

5.1.5 Verifying the requirements

Each of the system service components which we install has requirements which could include hardware, software, communication, database, or security. The Information Center not only contains the documentation for installing Virtualization Engine and the Virtualization Engine system service components, it also lists the requirements for each of the components which we install.

Before we began our installation, we reviewed the Information Center Virtualization Engine requirements as well as the Virtualization Engine console, IBM Director, and Tivoli Provisioning Manager and verified the requirements to make sure that they were met. Let us discuss these requirements in more detail.

Hardware requirements

Each product requires processor capacity, memory (physical, virtual, and swap), disk space (permanent and temporary), and potentially display characteristics, such as 256 color display. We verified that the total of the combined product requirements which we installed on a single xLinux are met because we installed multiple products on one system from the planning advisor Deploy tab to calculate these requirements. We selected the Portal and the management servers and added up the values. As shown in Figure 5-4 on page 92, the Virtualization Engine console requires 1.4 GHz microprocessor, 1 GB of memory, and 600 MB of disk space. IBM Director (shown by clicking the management server) requires 750+ Mhz Pentium® microprocessor, 768 MB memory (which includes the memory for the IBM Director bridge), and 316 MB disk space.

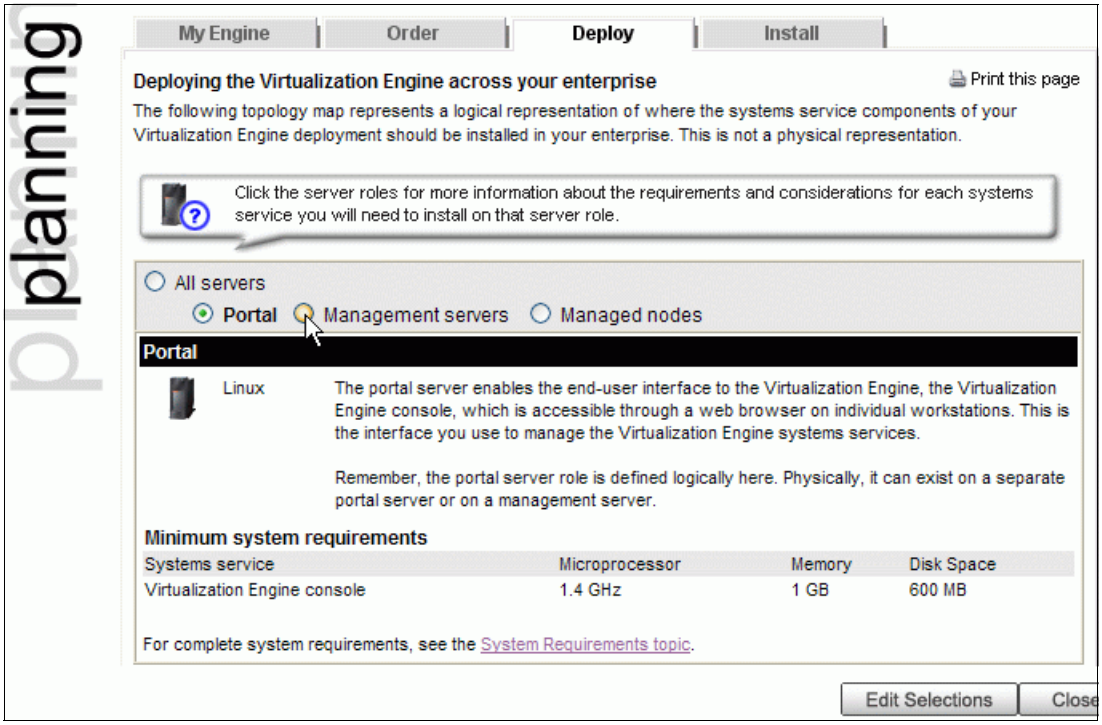
planning
My Engine
Order
Deploy
Install
Deploying the Virtualization Engine across your enterprise
Print this page
The following topology map represents a logical representation of where the systems service components of your Virtualization Engine deployment should be installed in your enterprise. This is not a physical representation.
Click the server roles for more information about the requirements and considerations for each systems service you will need to install on that server role.
All servers
Portal
Management servers
Managed nodes
Portal
Linux
The portal server enables the end-user interface to the Virtualization Engine, the Virtualization Engine console, which is accessible through a web browser on individual workstations. This is the interface you use to manage the Virtualization Engine systems services.
Remember, the portal server role is defined logically here. Physically, it can exist on a separate portal server or on a management server.
Minimum system requirements
Systems service
Microprocessor
Memory
Disk Space
Virtualization Engine console
1.4 GHz
1 GB
600 MB
For complete system requirements, see the System Requirements topic.
Edit Selections
Close

Figure 5-4 The planning advisor Deploy tab

Tip: When installing multiple products on one machine, it is best to verify the requirements for all products before installing the first one. For example, if multiple products are getting installed by different people, a conflict can arise for the ports that they are using or lack of space in a particular directory. The other situation is where there might be a common software prerequisite, such as DB2, which can require a higher fix pack level.

Planning for the best sequence of installation will eliminate potential problems.

The installation path of the prerequisite products is either on the /opt or /usr directories. It is recommended that these directories reside on a logical volume manager (LVM) or some other expandable file system. For example, when installing Tivoli Provisioning Manager, the prerequisite software products (WebSphere Application Server V5.1 and DB2 UDB Workgroup Server 8.1) get installed in /opt, while the prerequisite software product IBM Tivoli Directory Server gets installed in /usr.

Tip: Some of the installation paths for the prerequisite products cannot be changed using the installation wizards. Therefore, make sure there is enough disk space allocated to these directories.

Software requirements

For some of the products, there will be software requirements, such as additional rpms, device driver or fix packs. Additionally, when installing prerequisite software, such as the IBM Tivoli Directory Server, you might need to remove or migrate any currently installed Lightweight Directory Access Protocol rpms (LDAP servers, such as NSS LDAP from this system). Make sure that no other program or service are using these before removing them. We will detail this information as we discuss our installation in this chapter.

Network requirements

Determine and document what network protocols are supported for the different components (which are listed in the product documentation in the Information Center). Additionally, identify and document the ports that the products will use for communication. Default port numbers can have conflicts when installing new products. It is recommended to determine which protocols and ports are currently being used to potentially avoid a conflict. This can be accomplished by using the worksheets available with each of the Virtualization Engine components provided in the associated PDF in the Information Center.

Tip: It should be a good idea to list before any installation the port numbers already in use in your environments, to understand the new port numbers you need to implement when installing the Virtualization Engine, check that there is no contradiction and maybe change accordingly.

Security requirements

There are additional security requirements that can be established for the Virtualization Engine components. SSL and encryption are among these, although we chose not to use these for our environment. Instead, the security that we chose is based on authentication and authorization provided by the Tivoli Directory Server.

5.2 The installation process

In this section, we detail the installation process of the Virtualization Engine console, the IBM Director, and the Tivoli Provisioning Manager components.

5.2.1 Installing the Virtualization Engine console component

After verifying that the requirements have been met for the Virtualization Engine suite of products we are going to install, we proceeded with the installation of the Virtualization Engine console. These steps include:

1. Installing and configuring the Tivoli Directory Server.
2. Installing the Virtualization Engine console code.
3. Installing the Virtualization Engine console bridge to IBM Director.
4. Verifying the installation.

Tips:

- We found that clicking the planning advisor hyperlinks to get to the correct section in the installation process in the Information Center was very helpful.
- Additionally, having the associated PDF for each of the components was also extremely useful in order to fill in the worksheets and to check off when a step had been completed.
- For example, in the Virtualization Engine console installation, we need to define the LTPA key and specify the directory where it resides. When installing the IBM Director, we needed to import the LTPA key and used the information recorded in the Virtualization Engine console manual worksheet to obtain the name of the file and the directory in which it resides.

We printed the Virtualization Engine console V1R1 PDF that is available from the Information Center before we started this step to allow us to fill in the worksheets beforehand. This helped

us to identify many of the userids, TCP/IP, LDAP, and IBM Director definitions for successful installation and use of the Virtualization Engine console.

We completed the Virtualization Engine console worksheet items with the following answers:

Table 5-1 Completed Virtualization Engine console worksheet

Installation Screen Label	Answer
LDAP server host	velin12.itso.ibm.com
LDAP server port	389
DN used for directory updates	cn=root
DN of user and group tree	o=IBM,c=US
User naming attribute	uid
Relative location of users	cn=users
user object class	inetOrgPerson
Group naming attribute	cn
Relative location of groups	cn=groups
Group object class	groupOfUniqueNames
Group membership attribute	uniqueMember
DN for user and group lookup	uid=iscadmin,cn=users,o=IBM,c=US
SSO domain name	itso.ibm.com
ISC host name	velin12.itso.ibm.com
Starting port	8421
Portal administrator DN	uid=iscadmin,cn=users,o=IBM,c=US
Short user name	iscadmin
Portal administrator group DN	cn=iscadmins,cn=groups,o=IBM,c=US
Short group name	iscadmins

Step 1: Installing and configuring Tivoli Directory Server

For this step, we followed the steps in the Information Center under the section *Installing and configuring the IBM Directory Server 5.1 on Linux* within the Virtualization Engine console bullet with the following exceptions or clarifications.

Attention: In the following note that *Step n* (where *n* is a number) refers to the Information Center step and *Number n* (where *n* is a number) refers to the associated number under that step in the Information Center

To install and configure Tivoli Directory Server, do the following:

1. Prepare the system for installation, Number 1.

   The Information Center instructed us to remove all LDAP servers via the **rpm** command. However, we did not remove all installed LDAP servers because other RPMs had dependencies on them and were not related to this LDAP install.

The RPMs that we removed were:

- `nss_ldap-199-186`
- `openldap2-2.1.4-183`
- `pam_ldap-150-79`

We continued with the Information Center steps exactly until we reached the next step.

2. Install IBM Directory Server: Number 8.

   Although it is not articulated in the Information Center, at this point, we chose TYPICAL installation (instead of a custom installation). This option installs the following automatically:

   - IBM Tivoli Directory Server 5.1 into /usr/ldap
   - Its prerequisite products, DB2 V8.1 into /opt/IBM/db2/V8.1,
   - GSKit into /usr/local/ibm/gsk5
   - The embedded version of WebSphere Application Server - Express V5.0 into /usr/ldap/appsrv

   Note: By selecting TYPICAL as the setup type, you cannot overwrite the installation path directories for DB2, GSKit and the embedded version of WebSphere Application Server - Express 5.0.

   We continued with the Information Center steps exactly until we reached the next step.

3. Install IBM Directory Server: Number 13.

   Make sure you wait until the window appears for the Configuration Tool before clicking **Finish** because it might cause the Configuration Tool to not start automatically. If it does not appear, you can start it manually with the **`ldapxcfg`** command

   We continued with the Information Center steps exactly until we reached the next step.

4. Configure the IBM Directory Services 5.1 Server, Number 2g.

   We used the IBM Directory Server Configuration Tool to create the database, the DB2 instance, the DB2 instance node, start the database manager and enable loopback.

   We continued with the Information Center steps exactly until we reached the next step.

5. Configuring Virtualization Engine console LDAP settings using the IBM Directory Server Web Administration Tool, Number 7n.

   We started another Web browser to authenticate iscadmin.

After we installed the IBM Tivoli Directory Server and configured the Virtualization Engine console LDAP settings, the next step is to install the Virtualization Engine console code itself.

Tip: We entered many definitions while configuring the Virtualization Engine console LDAP settings. We used the Virtualization Engine console LDAP worksheet that is provided under the *Installing the Virtualization Engine console* section to simplify the next step.

Step 2: Installing the Virtualization Engine console code

For this step, we followed the directions from the Information Center under the section *Installing the Virtualization Engine console*.

We used the Virtualization Engine installation wizard to install the Virtualization Engine console. We chose the default ports and used the information from the Tivoli Directory Server install to fill out the LDAP panels as listed in Table 5-1 on page 94.

Because we are installing the Virtualization Engine console before the console bridge, we chose the option to export the LTPA key file (LTPAKeys.dat) using the **`export`** command. Take note of this file and its location, because you specify this file when installing the Virtualization Engine console bridge to IBM Director.

Step 3: Installing the Virtualization Engine console bridge

To connect the Virtualization Engine console to its management source, the IBM Director, we needed to define a bridge and configured it on the IBM Director system (which in our case is the same system as the Virtualization Engine console). The bridge contains Web Services that run on the management source system and allow a connection to the Virtualization Engine console Health Center and the IBM Director.

We used the Virtualization Engine installation wizard (the same one that we used to install the Virtualization Engine console) and selected the Virtualization Engine console bridge. We followed the directions in the Information Center under the section *Virtualization Engine bridge for Director Multiplatform or IBM Director 4.2 installation walkthrough* within the Virtualization Engine console section using the defaults and completing the worksheet provided, as shown on Table 5-2.

Table 5-2 The completed Virtualization Engine console bridge worksheet

Description	Our Value
WebSphere Application Server instance port number	12000
LTPA Timeout	120
LTPA Key File	/opt/IBM/VE/LTPAKeys.dat
Single Signon Domain Name	itso.ibm.com

Step 4: Verifying

Following the order of these procedures, we cannot verify at this time. If the IBM Director was installed before the Virtualization Engine console, then you could. Because the Virtualization Engine console (and the bridge) has been installed, we cannot verify until we have the IBM Director installed, which is in the next step.

5.2.2 Installing the IBM Director component

We went back to the planning advisor and selected Installing IBM Director Server. The installation flowchart appeared, and we selected the first prerequisites: Complete the IBM Director planning checklists. This opened a browser window for Information Center Planning for IBM Director. We selected **Checklist: Linux planning for Director Multiplatform** and reviewed the tasks that were identified.

Because hardware, software, and communication protocols and ports as well as the IBM Director Agents have already been met and identified, the only checklist item open was to make sure our database requirements have been met. Therefore, we will summarize our checklist items and detail what we needed to do to meet the database requirements.

Step 1: Determining where to install the IBM Director components

We defined our environment previously using two management servers. We use the first xSeries Linux machine to configure the Virtualization Engine console, Virtualization Engine console bridge and the IBM Director. We defined our managed servers to be Windows 2003,

AIX 5.3 and xSeries Linux SuSE Linux Enterprise 8. The managed servers would need to have the IBM Director Agent installed on them.

Step 2: Ensure that the system runs a supported operating system

Our management servers were running xSeries for Linux using SuSE Linux Enterprise Server 8, a supported operating environment and our managed servers were running Windows 2003, AIX 5.3 and xSeries Linux SuSE Linux Enterprise 8 (all supported operating systems).

Step 3: Ensure hardware and software requirements have been met

IBM Directory Server requires 750+ Mhz Pentium microprocessor, 768 MB memory (which includes the memory for the IBM Director bridge) and 316 MB disk space and a Display device of at least 256 colors.

The IBM Director Console require an additional 128 MB memory (RAM) and 168 BM disk space as we are running them both on the same server.

The IBM Director Agent has requirements which are listed in Table 5-3:

Table 5-3 IBM Director Agent hardware requirements

System Service	Microprocessor	Memory	Disk Space
IBM Director Agent (AIX)	0-2% of application CPU[1]	128MB	43MB
IBM Director Agent (Linux)	0-2% of application CPU[1]	128MB	88-109MB
IBM Director Agent (Windows)	0-2% of application CPU[1]	128MB	109MB

1. The IBM Director Agent runs on any IBM processor that is supported by the current operating systems. Some monitor functions can use approximately 1% to 2% of the processing power that is used by the application if no monitors are configured.

There are no additional software requirements for xSeries Linux SuSE Linux Enterprise 8.

There were specific device driver requirements for Linux on Xseries. We needed to verify that we had LM78 and SMBus device drivers installed to ensure that certain IBM Director functions and tasks work properly, which we did.

Step 4: Reviewing the communication protocols and identify ports

We verified that we had TCP/IP V4.0. We did not require SNA, NetBIOS or IPX. The ports that we used were the default ports as documented in the Virtualization Engine IBM Director V1R1 PDF (or available in the Information Center under IBM Director Multiplatform → Requirements → ports).

Step 5: Ensure that the database meets the requirements

We chose to use the same local DB2 database as the Virtualization Engine console because both products would be running on the same system. However, IBM Director requires DB2 Universal Database 8.1, Fix Pack 5 be applied. In our case, we did the following:

- We stopped all applications which were using the DB2 which in our case was only the Virtualization Engine console, then we stopped DB2. We went into the DB2 fix pack directory where we downloaded and unzipped the DB2 fix pack code and issued the **`installFixPak`** command which started a wizard which allowed us to install the fix pack to the DB2 installation.

- We updated the DB2 instance by running **db2iupdt** and reapplied the licence with the **db2licm** command. We then restarted DB2 and the DB2 database manager, as documented in the release.txt file in the fix pack.

Step 6: Understanding security and setting up the appropriate privileges

We decided not to use SSL as this is a test environment. To use IBM Director, a user must have an operating-system account on the management server or the domain and be authorized to use IBM Director. There are two user groups that get installed with IBM Director: Administrator group called *diradmin*, and Super-user group, called *dirsuper*. After we install the IBM Director, we will need to add to this group, as described in “Step 8: Installing IBM Director code” on page 98.

Step 7: Determining which IBM Director Agent features have been identified

We installed all of the available IBM Director Agent features.

Step 8: Installing IBM Director code

We used the Virtualization Engine installation wizard (the same one that we used to install the Virtualization Engine console) and selected the IBM Director option to install IBM Director. We then selected all of the IBM Director components to install, as shown in Figure 5-5.

Figure 5-5 IBM Director Server selections

After the installation, as documented in the Information Center under IBM Director Multiplatform → Configuring → Configuring the database on Linux, we did the following:

- Configured the database.

  When configuring the database, we chose to use the same DB2 instance as the Virtualization Engine console, db2inst1, and we named the IBM Director’s database idmdb.

  We ensured that the server for IBM Director was stopped. To stop the server, enter the following command:

  ```
  /opt/IBM/VE/director/bin/twgstop
  ```

We created a /etc/TWGserver/setup_env file and added the following statement to the file:

```
/home/db2inst1/sqllib/db2profile
```

In this statement, `home/db2inst1` is the directory where DB2 Universal Database is installed. This statement sets up the DB2 Universal Database environment.

We set the setup_env file attributes to read-execute.

We configured the database for use with IBM Director using the following command:

```
/opt/IBM/VE/director/bin/cfgdb
```

- Authorized users to use IBM Director.

  To use IBM Director, a user must have an operating-system account on the management server or the domain and be authorized to use IBM Director. There are two user groups that get installed with IBM Director: Administrator group called `diradmin`, and Super-user group, called `dirsuper`. We added the user `director` and the user `root` to the `dirsuper` group because this group can define the privileges to the administrator group as well as create and edit individual user accounts with the following commands:

  ```
  usermod -G root,pkcs11,other,dirsuper root
  ```

- Started IBM Director using the following command:

  ```
  /opt/IBM/VE/director/bin/twgstart
  ```

After we installed and configured the IBM Director, we started the IBM Director Console, which gives us access to the graphical user interface to perform various IBM Director tasks using the following command:

```
/opt/IBM/VE/director/bin/twgcon
```

5.2.3 Installing the Tivoli Provisioning Manager component

There are several topologies that are supported by Tivoli Provisioning Manager and the supporting software, such as the database server and the Web server. These topologies include:

- A one-node topology where Tivoli Provisioning Manager and the supporting software are installed on a single machine. This includes Tivoli Provisioning Manager, the database server, WebSphere Application Server, and the LDAP server.
- A two-node topology where the LDAP server is installed on a separate machine, while Tivoli Provisioning Manager, the database server, and WebSphere Application Server, are installed on another.
- A three-node topology where the LDAP server is installed on a separate machine, the database server is installed on a second node and Tivoli Provisioning Manager and WebSphere Application Server, are installed on a third node.

Important: In this exercise, we installed Tivoli Provisioning Manager in a one-node topology to reduce the complexity of our environment.

Because Tivoli Provisioning Manager requires Fix Pack V2.1.0.1 to be supported on SuSE 8 on xSeries, after installing Tivoli Provisioning Manager 2.1, we installed Fix Pack 1 immediately as described in “Applying fix packs to the Tivoli Provisioning Manager” on page 110. Additionally, the following software was required:

- WebSphere Application Server 5.1 plus fixes located on the Tivoli Provisioning Manager Generic Fixes CD
- DB2 8.1 plus Fix Pack 3

- IBM Tivoli Directory Server 5.2 client and server, plus fixes located on the Tivoli Provisioning Manager Generic Fixes CD

After verifying that the requirements have been met as documented in *Systems Provisioning with IBM Tivoli Provisioning Manager Planning,* BOEE-EUSP-01, we proceeded with the installation of Tivoli Provisioning Manager using *TPM Installation Guide for Linux*, GC32-1616.

Tip: The *TPM Installation Guide for Linux*, GC32-1616 is a useful manual for planning the Tivoli Provisioning Manager installation and identifying the supported hardware and software, as well as describing the supported Tivoli Provisioning Manager topologies and needed CDs.

We followed the steps outlined in the *TPM Installation Guide for Linux*, GC32-1616, for installing and configuring the system prerequisites.

Step 1: Installing and configuring the system prerequisites

To install and configure the system prerequisites, we performed the steps listed in this section.

Step 1.1: Installing and configuring the WebSphere Application Server prerequisites

You have to install and configure utilities and user IDs before you install Tivoli Provisioning Manager. In our configuration, the WebSphere Application Server is a prerequisite, for which the following steps need to be achieved. Before you install WebSphere embedded messaging, you must log in as root and perform the following tasks:

1. Create the file systems for WebSphere MQ.
   a. Before you install WebSphere MQ, create and mount a journalized file system called /var/mqm. Use a partition strategy with a separate volume for the WebSphere MQ data. This ensures that other system activity will not be affected if /var/mqm accumulates a large amount of WebSphere MQ data.
   b. Ensure that the /var directory is large enough to handle the workload. The /var file system is used to store all the security logging information for the system, and is used to store the temporary files for email and printing. Therefore, it is critical that you maintain free space in /var for these operations. If you do not create a separate file system for messaging data, and /var fills up, all security logging will be stopped on the system until free space becomes available in /var, and you will not be able to email or print. To determine the size of the /var/mqm file system for a server installation, consider the following criteria:
      - The maximum number of messages in the system at one time
      - Contingency for message buildups, if there is a system problem
      - The average size of the message data, plus 500 bytes for the message header
      - The number of queues
      - The size of log files and error messages
      - The amount of SSL trace that is written to the /var/mqm/trace directory

   Allow 50 MB as a minimum for a WebSphere MQ server. You need less space in /var/mqm for a WebSphere MQ client. Typically this is 15 MB.

2. Create the file systems for the WebSphere MQ working data.
   a. Create a file system /var/mqm/log for your log data.

      If possible, store log files on a different physical volume from the WebSphere MQ queues (/var/mqm). This ensures data integrity in the case of a hardware failure.

   b. Create a file system /var/mqm/errors for your error files.

      Note: The size of the log file depends on the log settings that you use. The recommended size is set to support circular logging using the default settings. For further information about log sizes, see *WebSphere MQ System Administration Guide*, SC34-6068.

3. Set up the group IDs and user ID for WebSphere MQ.

   You must create the specified group IDs and user ID before you install WebSphere MQ.

   a. Create the mqm group and mqbrkrs group.
   b. Create the user mqm with mqm as the primary group. The mqm user ID owns the directories and files that contain the resources associated with WebSphere MQ.
   c. Add the mqm user to the mqbrkrs group.
   d. Add the user root to the mqm group and mqbrkrs group.
   e. Log out as root.
   f. Log in as mqm user that was just created.
   g. Log out as mqm.
   h. Log in as root to implement the change and make sure your changes were applied by running the appropriate command: **`id -a`**

Step 1.2: Installing the prerequisite software

Ensure that you have installed the following packages:

- Expect 5.34 is a tool for automating keystrokes and must be installed prior to installing Tivoli Provisioning Manager. It can be installed through the package installer for both Linux iSeries and Red Hat Advanced Server 2.1. The level we were running with was expect-5.3-192.
- Because we were not running Tivoli NetView® we did not require SNMP.
- Tivoli Provisioning Manager makes extensive use of the SSH, telnet and ftp protocols. Ensure these protocols are available on the Tivoli Provisioning Manager servers. Follow the instructions in this section after installing Tivoli Provisioning Manager. Do not configure SSH before installing Tivoli Provisioning Manager.
- libgcc: This Linux package is required by the Tivoli Directory Server client. The level needs to be at libgcc-3.0.401 or higher. We validated that we were at libgcc-3.2.2-54 with the following command:

  ```
  rpm -qa | grep libgcc
  ```

Step 2: Installing the Linux prerequisite software

We then continued with installing the software prerequisites on Linux Intel using the prerequisite software installer as described in the *TPM Installation Guide for Linux*,

GC32-1616. However, we needed to make the following alterations because we were running on SuSE 8:

1. Log on as root.
2. Insert the Tivoli Provisioning Manager and Intelligent ThinkDynamic Orchestrator Version Tivoli Provisioning Manager 2.1 Prerequisite Software Installer CD into the CD drive. Mount the CD-ROM drive, but do not change directory to the mount point.

Tip: Changing directories to the mount point locks the CD drive and prevents you from being able to swap CDs. You must unmount the CD-ROM before trying to eject the CD. Otherwise, the CD-ROM tray will lock, and you will be unable to switch to CDs of the other software prerequisites.

3. Enter the following command:

```
mount_point/setuplinux.bin
```

4. Select the language.

   A language selection panel for the wizard appears. Select which language in which you want the installation panels to display. Support for this language will also be installed for the prerequisite software products chosen. To install additional languages, use the language pack install process provide by each of the individual prerequisite software products. Click **OK**. The installer program verifies that the system prerequisites have been met.

5. On the Prerequisite software installer panel, click **Next**.
6. Specify a directory.

   On the next panel, specify a directory into which you will copy the Tivoli Software Installer application, and then click **Next**. The Tivoli Software Installer copies itself to the local hard drive to free up the CD drive. While the Tivoli Software Installer copies code to the hard drive, the installer screen might not be visible.

   The language selection panel for the wizard appears again. Click **OK**.

7. Select the products.

   Select which products to install on this machine. Choose from these options:

   – DB2 Universal Database 8.1: Installs either the DB2 server or client. Choose this option if you want to install only the database client or the server.
   – Tivoli Directory Server 5.2 with DB2 8.1: Installs the Tivoli Directory Server client along with DB2.

Tip: If you select DB2 Universal Database 8.1, you get a full DB2 install (db2das1 and db2fenc1, and so forth. You can come back later to select Tivoli Directory Server 5.2 with DB2 8.1. If you select Tivoli Directory Server 5.2 with DB2 8.1, you will not have a complete DB2 installation but one that is tailored for IBM Tivoli Directory Server.

Tip: To install the server, use the install wizard to install on an operating system other Red Hat Linux 2.1 or AIX 5.2.

8. When you have made your selection, click **Next**.

   Note: While the prerequisite software is being installed, the progress bar will only be updated after each software product has been installed.

9. If you selected to install DB2 Universal Database, the next panel allows you to specify whether to install the client or server. Choose from these options:
   - DB2 Workgroup Server with Fix Pack 3: This will install the DB2 Universal Database, Workgroup Server Unlimited Edition, Version 8.1 database server. Select this option if one of the following conditions applies:
     - You want to install your database on the same machine as Tivoli Provisioning Manager.
     - You want to install your database on a dedicated machine.
   - DB2 Administrator Client with Fix Pack 3: This will install the DB2 client. Select this option if you will be creating, or have already created, a remote database server and you are running the installer on your Tivoli Provisioning Manager machine. The DB2 client instance will be created for you.
10. Data base configuration.
    - On the next panel, you are prompted to enter information which will be used to configure the database and create your Tivoli Directory Server. Select either:
      - Server and Client
      - Client Only

      Because we are running with the local DB2, we did not need the client, so we selected Server and client.

      Attention: This step is not described explicitly in the documentation.

    - On the next panel, you are prompted to enter information which you will use to configure the database and to create your database instance. Complete the fields as follows:

Database user name	You can change the default user name of db2inst1 but do not leave this field blank.
Password	Enter the password for the database user name.
Confirm password	Enter the password again.

11. Install Tivoli Directory Server client.

    If you selected to install Tivoli Directory Server 5.2 with DB2 8.1, the next panel gives a message that only the client portion of the Tivoli Directory Server 5.2 with DB2 8.1 will be installed if you are running on Red Hat. This does not occur on SuSE. This installs the Version 5.2 client *only* and not the server.

    Before you continue, you are presented with the installation options that you have selected, including installation directories, port numbers, user names, and so on. Record this information. Some of the information will be needed during the Tivoli Provisioning Manager installation. The installation options are also saved in usr/ibm/pics/mwInfo.html

12. Click **Install**.
13. Provide the installation code location.

    You are prompted to enter or to browse to the location of the installation code for the software that you have selected. If you are installing the software selections from the local

CD drive, insert the first CD and specify the appropriate drive location. If you have copied the software CD images to a local or network drive, enter the path to the location of the images.

Depending on the software that you have selected, you are prompted for installation images in the following order:

a. DB2 Universal Database, Workgroup Server Unlimited Edition, Version 8.1 (see ATTENTION notice)

b. DB2 Universal Database Version 8.1 Fix Pack 3

c. Tivoli Directory Server Version 5.2

d. WebSphere Application Server 5.1

14. Click **Next** after you have selected each option.

Attention: The graphical installation process selects port number 50000 by default. However, this port is already assigned by the system. You can workaround this problem in one of two ways:

- Manually install DB2 without using the graphical installer. To follow this workaround, refer to the DB2 installation documentation.
- Modify the db2server.txt file so that the graphical install process functions properly:

  a. Start the graphical installer and follow the panels to the prompts for DB2 instance owner and the DB2 password. You will then be prompted for the directory location of the DB2 workgroup server package.

  b. Before you enter the directory location, you must correct the port number conflict:

     - Open /opt/IBM/db2/V8.1/pics/db2server.txt.
     - In the db2server.txt file, locate the following: `db2inst1.PORT_NUMBER = 50000`.
     - Change the port number from `50000` to `50001`.
     - Save the file.

  c. Return to the graphical installer and enter the location of the DB2 workgroup server package and continue the graphical install process.

Tips:

- Do not drill down to the location of the actual installation executable. Select the top-level directory of the software location.
- When prompted for the location of the DB2 install image, if you are using Oracle as your database server and you do not want to install the DB2 client or the DB2 server, click Cancel. This cancels the installation wizard.
- The DB2 Universal Database and WebSphere Application Server CDs have autorun enabled. Close the install panels for these applications when they are displayed.

15. Acknowledge the final step.

    Click **Finish**. You have now installed the prerequisite software required for Tivoli Provisioning Manager.

16. Add the DB2 license with the **`db2licm`** command as follows:

    a. Log in as root.

    b. Insert the DB2 CD into the CD drive. Mount the CD.

    c. Switch to the DB2_installdir/adm/ directory where the **`db2licm`** utility can be found.

d. Issue the following command to apply the license (for a permanent license):

```
db2licm mountpoint/db2/license/db2wsue.lic or db2licm -l
```

For more information about the **db2licm** command, refer to *DB2 Command Reference*, S20H-4645.

e. Issue the following command and replace the variable *filename* with the name of the licensing file:

```
db2licm filename.lic
```

For more information about the **db2licm** command, refer to *Command Reference* in the DB2 documentation.

You should receive a message similar to the following:

```
DBI1402I License added successfully.
DBI1426I This product is now licensed for use as specified in the License
Acceptance and License Information documents
```

Step 3: Configuring Tivoli Directory Server

We next configured the Tivoli Directory Server. After we configured the IBM Tivoli Directory Server, we returned to Chapter 9 of *TPM Installation Guide for Linux*, GC32-1616, to apply the cumulative fix 3 to WebSphere Application Server and the IBM WebSphere MQ V5.3 fix.

In Chapter 9, the first step is to apply the Tivoli Directory Server fix. You cannot do this unless you have configured the server, which is described in Chapter 10 of *TPM Installation Guide for Linux*, GC32-1616.

Attention: In the documentation version that we used, the order of the chapters is in error. You have to follow the instructions in Chapter 10 before those in Chapter 9. If you do Chapter 9 steps before Chapter 10 steps, the steps will fail. The documentation might be fixed when you read this redbook.

This section describes how to configure Tivoli Directory Server using the Configuration Utility. You can use the **ldapcfg** utility to configure Tivoli Directory Server. For more information, refer to the Tivoli Directory Server *Installation and Configuration* documentation, which is available at:

http://publib.boulder.ibm.com/tividd/td/IBMDirectoryServer5.2.html

To configure Tivoli Directory Server, do the following:

1. Log on as root.
2. From a command prompt, enter the **ldapxcfg** command to start the Tivoli Directory Server configuration utility.

Tip: Do not minimize the Configuration Tool window or the command prompt window that is displayed during initial configuration. Unpredictable results might occur.

Setting the Administrator Domain Name and password

To configure the Administrator Domain Name (DN), do the following:

1. In the task list, click **Administrator DN**.
2. Enter a valid domain name (or accept the default domain name of cn=root) in the Administrator DN field. The Tivoli Directory Server Administrator DN is the domain name that is used by the administrator of the directory. This administrator is the one user who has full access to all data in the directory. The default domain name is cn=root. Domain

names are not case sensitive. If you are unfamiliar with X.500 format or if for any other reason you do not want to define a new domain name, accept the default.

3. Type the password for the Administrator DN in the Administrator Password field. You must define a password. Passwords are case-sensitive. Record the password for future reference.
4. Retype the password in the Confirm password field. Click **OK**.

Note: Double byte character set (DBCS) characters in the password are not supported.

Configuring the database

When you configure the database, the configuration tool adds information about the database that will be used to store directory data to the configuration file ibmslapd.conf. Before you configure the database ensure you complete the following tasks:

1. Stop the directory server before you configure the database.
2. Ensure that the environment variable DB2COMM is not set.
3. Create a user ID for the user who will own the DB2 database. The user ID that you specify will own the database instance where the DB2 database will exist, and the DB2 instance will be in the home directory of the user. The user ID can be no longer than 8 characters. In addition:
   a. The user must have a home directory and must be the owner of the home directory.
   b. The group ownership of the home directory of the user should be the DB2 group created when DB2 was installed. The user root must be a member of the primary group of the user. If root is not a member of this group, add root as a member of the group.
   c. The password of the user must be set correctly and ready to use. For example, the password cannot be expired or waiting for a first-time validation of any kind. The best way to verify that the password is correctly set is to telnet to the same computer and successfully log in with that user ID and password.
   d. When configuring the database, it is not necessary, but customary, to specify the home directory of the user ID as the database location. However, if you specify some other location, the home directory of the user still must have 3 to 4 MB of space available. This is because DB2 creates links and adds files into the home directory of the instance owner (that is, the user) even though the database itself is elsewhere. If you do not have enough space in the home directory, you can either create enough space or specify another directory as the home directory.

      To configure the database:

      i. In the Configuration Tool, click **Configure database** in the task list.
      ii. The Configuration Tool attempts to determine whether you already have a database. If you have a database already configured (that is, the information for the database is in the configuration file), the Configuration Tool prompts you for information about what you want to do. For example, if the database is configured but cannot be found on the system, you might choose to create a database using the name specified in the configuration file. Use the information shown in the windows that are displayed to configure the database.

      Note: The windows that are displayed in the remainder of these steps depends on whether you already have a database.

iii. If a user ID and password are requested: a. Type a user ID in the **User ID** field. This user ID must already exist before you can configure the database. b. Type a password for the user in the Password field. Passwords are case-sensitive.

iv. If the database name is requested, you need to enter the name that you want to give the DB2 database. The name can be from 1 to 8 characters long. The database will be created in an instance with the same name as the user ID. After you enter the name, click **Next**.

> **Note:** If you want a different database instance name, you must use the **`ldapcfg`** command with the **`-t`** option to configure the database.

v. If the database location is requested, you need to enter the location for the database in the Database location field. For non-Windows platforms, the location must be a directory name, such as /home/ldapdb2. When you have entered the location, click **Next**.

> **Note:** Be sure that you have at least 80 MB of free hard disk space in the location that you specify and that additional disk space is available to accommodate growth as new entries are added to the directory.

vi. If a character set selection is requested, select the type of database that you want to create. You can create a UCS Transformation Format (UTF-8) database, in which LDAP clients can store UTF-8 character data, or a local code page database, which is a database in the local code page. When you have select the type of database, click **Next**.

> **Note:** If you want to use language tags, the database must be a UTF-8 database.

vii. Information is displayed about the configuration options that you have specified. To return to an earlier window and change information, click **Back**. To begin configuration, click **Finish**. The completion window is displayed.

viii. Click **Close**.

Step 4: Applying mandatory patches to WebSphere Application Server and Tivoli Directory Server

This step explains how to apply the mandatory patches for WebSphere Application Server and Tivoli Directory Server.

Applying the WebSphere Application Server 5.1 Cumulative Fix 3

To apply the WebSphere Application Server 5.1 Cumulative Fix 3:

1. Log on as root.
2. Stop the WebSphere Application Server.
3. Switch to the WebSphere Application Server installation directory and create a sub-directory called /update/was510_cf3.

4. Copy the Cumulative Fix 3 file from the WAS/CF3 directory on the CD to the *WAS_installdir*/update/was510_cf3 directory you created. Depending on your platform the file name is as follows:
   - On Linux Intel: was510_cf3_linux.zip
   - On Linux for iSeries: was510_cf3_linuxppc.zip
5. Uncompress the Cumulative Fix 3 package.
6. Run the UpdateWizard with the **updateWizard.sh** command. During installation, ensure that the Install fix packs option is selected.

   Note: On some platforms, the UpdateWizard interface to the update installer program does not recognize the existing WebSphere Application Server product. This problem is caused by a limitation in the InstallShield for MultiPlatforms (ISMP) program that the update installer program uses. To workaround the problem, click **Specify product information** and type the fully qualified installation root directory for the existing product in the Installation directory field of the UpdateWizard panel.

7. After installing the fix, restart WebSphere Application Server.

Applying the IBM WebSphere MQ V5.3 fix

To apply the WebSphere MQ fix:

1. Logon as root.
2. Copy the file IC38409.unix.tar.Z file from the CD located in the directory: /WAS/IC38409 into a temporary directory on your hard drive.
3. Uncompress the file using the following command:

   ```
   uncompress IC38409.unix.tar.Z
   ```

4. Untar the IC38409.unix.tar file into a temporary directory on your hard drive using the command:

   ```
   tar -xvf IC38409.unix.tar
   ```

5. Stop all running queue managers and channel listeners.
6. Write down the properties of the existing com.ibm.mqjms.jar and com.ibm.mq.jar files (permissions, ownership).
7. Make a backup copy of the existing files.
8. Copy the new versions on top of the existing com.ibm.mqjms.jar and com.ibm.mq.jar files. Check that the new versions of files have retained the same properties as the old versions. If required, correct any modified properties.
9. Restart the queue managers and channel listeners.

Step 5: Installing Tivoli Provisioning Manager code

After we installed the prerequisite software, we then installed the Tivoli Provisioning Manager using the graphical installer as described in 5.2.3, “Installing the Tivoli Provisioning Manager component” on page 99.

Attention: During this project, we had an issue with the Host Port definition. The problem only occurred on the SuSE Linux system, not the Red Hat Linux which we tested in another installation. During the installation of Tivoli Provisioning Manager with the SuSe Linux code that we used, the Database Server Configuration panel is displayed and prompts you for the information required to configure your database server. One of the items is the Fully Qualified Host Name or IP Address of the IBM DB2 Universal Database server as well as the Host port.

Although we changed the port number to 50001 when installing DB2 as described “Step 2: Installing the Linux prerequisite software” on page 101, the port that the Tivoli Provisioning Manager installer uses is the ldapdb2svc port of 3700 which we obtained from the /etc/services file.

Step 6: Post-installation steps for Tivoli Provisioning Manager

After you have installed Tivoli Provisioning Manager, you must create a data center model (DCM) which Tivoli Provisioning Manager uses as described in 3.2.1, “The data center model” on page 41. When you have created the DCM, you have to then import it to Tivoli Provisioning Manager. This XML file is provided with the Tivoli Provisioning Manager installation and is called datacentermodel.xml. This file provides the information that Tivoli Provisioning Manager uses to load drivers, set up devices, and create the data center infrastructure, such as the systems and their components, represented by objects.

5.3 Fix pack installation considerations

You install a fix pack by using a fix pack installation wizard, which is provided as an executable file within the archived fix pack that you download. The fix pack installation process is similar to the installation process for the Virtualization Engine itself.

To ensure that you have the latest Virtualization Engine fix pack, you need to check the Virtualization Engine services fixes and updates Web site and download the appropriate fix packs for the common runtime and your systems services. You can find the Virtualization Engine systems services fixes at:

`http://techsupport.services.ibm.com/server/VirtualizationEngine`

When installing fix packs, whether it is a Virtualization Engine fix pack for a system service or a fix pack for a product which a system service is dependent on, there are some items to consider. For more information about fix pack installation, see the Information Center under Virtualization Engine System services fix pack considerations.

Tip: Before you install a fix pack, always review the readme file for the most up-to-date information about the content of the fix pack, as well as any special notices or instructions that might apply.

Tip: Due to a limitation with the Virtualization Engine installation process, you can only install a fix pack for the Virtualization Engine console on Linux if the directory location for the console does not contain any spaces. For example, if the directory name is /VE Console versus /VEConsole, the fix pack installation will fail. This limitation does not apply to other supported operating systems nor does it apply to the console bridges or other systems services.

5.3.1 Applying fix packs to the Virtualization Engine console

The Virtualization Engine console and the console bridges have dependencies on components of the Virtualization Engine. For example, the console bridges have dependencies on WebSphere Application Server.

Because the Virtualization Engine installation wizard automatically stops and starts any WebSphere Application Server, you might want to install any WebSphere Application Server fix packs before installing the fix packs for the console bridges. For example, to install WebSphere Application Server fix pack 5.1.1 or later, you need to stop the WebSphere Application Server that the console bridges use. When you stop the console bridges, the WebSphere Application Server is also stopped. After you apply the needed fix pack for WebSphere Application Server, you can install any Virtualization Engine fix packs that apply to the console bridges.

5.3.2 Applying fix packs to the IBM Director

When you install a Virtualization Engine fix pack for IBM Director, you should understand the implications for uninstalling IBM Director after a fix pack is applied. If you install a fix pack for Linux, you cannot uninstall fix packs. To remove a fix pack for IBM Director on AIX, Linux, NetWare, or Windows, you must uninstall IBM Director completely and then reinstall the version that you want.

A full description of the prerequisites for installing and installing fix packs on IBM Director are available in the Information Center under IBM Director Multiplatform → Installing → Fix pack considerations.

Tip: The product levels and the fix levels of the IBM Director Server, Console, and Agent are very dependent each other. Be sure they are all synchronized as requested by IBM.

5.3.3 Applying fix packs to the Tivoli Provisioning Manager

Fix Pack V2.1.0.1 is required to support Tivoli Provisioning Manager on SuSE 8 on xSeries, therefore, before using the product we needed to install this fix pack. We followed the instructions in the readme file for the fix pack which were:

1. Log in to the Tivoli Intelligent ThinkDynamic Orchestrator server or Tivoli Provisioning Manager server as root.
2. Ensure that the database server and the directory server are both running.
3. Export the JAVA_HOME environment variable to *WAS_HOME*/java where *WAS_HOME* is the location of the WebSphere Application Server, using the following commands:

   ```
   echo $JAVA_HOME
   export JAVA_HOME=/opt/WebSphere/AppServer/java
   ```

4. Ensure that the DISPLAY variable is set and confirm that the host name of the server is correct using the following commands:

   ```
   echo $DISPLAY
   echo $HOSTNAME
   ```

5. Change directory to *update_installer_dir*/installer where *update_installer_dir* is the directory in which you unzipped the fix pack zipped file.
6. From /*update_installer_dir*/installer, run **updateWizard.sh** to launch the installation wizard.
7. Select the correct language and click **OK**.
8. Click **Next** on the Welcome page.

9. Specify the installation root directory.

> **Tip:** The installation directory might not be specified here because on some platforms the graphical update installer program does not recognize the existing Tivoli Intelligent ThinkDynamic Orchestrator or Tivoli Provisioning Manager installation. To workaround this limitation:
>
> - In the Installation Directory box, if the installation directory is not specified, select **Specify product information**.
> - Browse and select the appropriate installation root directory. For example: /opt/IBM/tivoli/thinkcontrol.

10. Click **Next** until you are prompted to select an option.
11. Select **Install fix packs** and click **Next**.
12. The fix pack directory screen displays the default directory where the fix pack is installed. Verify that the fix pack is in this directory and click **Next**.
13. On the fix pack selection screen, select the correct fix pack, and click **Next**.
14. Review the information about the review screen and then click **Next**. The summary screen is displayed.
15. On the summary screen, **Finish**. The system completes the installation.

> **Tip:** The installations produces many log entries which are displayed in the background display window. These entries are also logged in /usr/IBM/tivoli/thinkcontrol/logs/update.

To complete the fix pack installation after the fix pack has been installed successfully, there are two post-installation tasks which you need to perform:

1. Run the update_xmi-sh script as follows:
   a. Log in as user root.
   b. Download and untar the UPDATE_XMI.TAR package from the Web site, from which you have downloaded Tivoli Provisioning Manager Fix Pack (2.1.0.1-TIV-TPM-FP01) and save it to your temporary directory.
   c. Open a Bash window and run the update_xmi.sh script, from the directory in which you have saved it.
2. Run the postinstall.sh script.
   a. Log in as user tiodamin.
   b. Ensure that the WebSphere Application Server is stopped.
   c. To deploy the EAR (expanded archive) files and database schema change, follow these instructions:
      i. Change the directory to *update_installer_dir*/installer/utils where *update_installer_dir* is the directory in which you unzipped the fix pack zipped file.
      ii. From *update_installer_dir*\installer, run the **`postinstall.sh`** script and replace the variable *was_user* with your WebSphere Application Server user name, and *was_pwd* with your WebSphere Application Server password. The post-installation operation prints the output to the console.

> **Tip:** This command takes a while to complete because it updates the database, installs updated workflows, and installs drivers. Be patient!

Tip: To redirect the output to a log file, run the following command and replace the variable *was_user* with your WebSphere Application Server user name and *was_pwd* with your WebSphere Application Server password:

```
postinstall.sh was_user was_pwd > log.txt
```

Refer to the instructions in the readme file of the fix pack for further information about the post-installation tasks for the latest information.

5.4 Post-installation step: configuring the data center model

Part of the post-installation steps instructs you to edit the data center model (DCM) XML file and change the values of the following:

- @VLANID@
- @MACADDR@
- @NICIF@
- @MACHINEMODEL@
- @TIOPWD@

The post-installation steps then instruct you to import the file into Tivoli Provisioning Manager. This section discusses how we changed this file for our Linux environment.

5.4.1 The DCM that is provided with Tivoli Provisioning Manager installation

Example 5-1 shows the data center model that is provided with the Tivoli Provisioning Manager installation. The file name is datacentermodel.xml and is located in $TIO_HOME/xml directory.

As you can see, the server name and the IP address are completed automatically with the Linux image information. However, the default data center model assumes the Linux is a Red Hat Linux. Additionally, there are several definitions that these definitions are dependent upon that are missing from this file.

Example 5-1 The DCM provided with the Tivoli Provisioning Manager

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- edited with XMLSPY v5 rel. 4 U (http://www.xmlspy.com) by IBM Corporation (IBM
Corporation) -->
<!DOCTYPE datacenter SYSTEM "xmlimport.dtd">
<datacenter>
        <spare-pool name="Manage From Linux" os-type="Red Hat Linux" vlan="@VLANID@"
fabric="Default Fabric">
<server name="velin11">
<nic name="@NICIF@" macaddress="@MACADDR@" netboot-enabled="false" managed="true">
<network-interface name="@NICIF@" ipaddress="10.1.1.176"/>
</nic>
<property name="platform" value="@MACHINEMODEL@"/>
<sap name="scp client" port="22" host="false" app-protocol="SCP" protocol-type="ipv4">
<credentials search-key="primary" is-default="true">
<password-credentials username="root"/>
</credentials>
</sap>
<sap name="ssh service" port="22" host="true" app-protocol="SSH" is-device-model="SSH
Service Access Point" protocol-type="ipv4">
<default-sap operation-type="execute-command"/>
```

```
<credentials search-key="secondary" is-default="false">
<password-credentials username="tioadmin" password="@TIOPWD@"/>
</credentials>
<credentials search-key="primary" is-default="true">
<rsa-credentials username="tioadmin"/>
</credentials>
</sap>
<sap name="snmp-set" port="161" host="true" app-protocol="SNMP" is-device-model="SNMP V1
Service Access Point" protocol-type="ipv4" context="write">
<default-sap operation-type="set-attribute"/>
<credentials search-key="primary" is-default="true">
<snmp-credentials community="public"/>
</credentials>
</sap>
<sap name="snmp-get" port="161" host="true" app-protocol="SNMP" is-device-model="SNMP V1
Service Access Point" protocol-type="ipv4" context="read">
<default-sap operation-type="get-attribute"/>
<credentials search-key="primary" is-default="true">
<snmp-credentials community="public"/>
</credentials>
</sap>
<software-association software=" RedHat Linux" state="running"/>
<software-association software="TIO" state="running"/>
</server>
</spare-pool>
<kanaha-config>
<dcm-object id="0">
<property component="DEPLOYMENT_ENGINE" name="PackageIntermediateHost" value="velin11"/>
<property component="DEPLOYMENT_ENGINE" name="PackageRepositoryHost" value="velin11"/>
<property component="DEPLOYMENT_ENGINE" name="PackageWebHost" value="velin11"/>
<property component="DEPLOYMENT_ENGINE" name="SoftwareRepositoryHost" value="velin11"/>
</dcm-object>
</kanaha-config>
</datacenter>
```

5.4.2 Our DCM XML file

We added the top section to the XML file which defines the switch fabric, the VLANS, and the subnet name, as shown in Example 5-2.

Example 5-2 The DCM XML file after customization

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- edited with XMLSPY v5 rel. 4 U (http://www.xmlspy.com) by IBM Corporation (IBM
Corporation) -->
<!DOCTYPE datacenter SYSTEM "xmlimport.dtd">
<datacenter>
        <switch-fabric name="ITSO_Fabric" />
        <switch name="ITSO_Switch" locale="en_US" fabric="ITSO_Fabric" failed="false"
in-maintenance="false">
<switch-module>
                        <switch-port failed="false" managed="false" netboot-enabled="false"
vlan="1" port-number="1" enabled="true" layer1-interface-type="unknown" />
                        <switch-port failed="false" managed="false" netboot-enabled="false"
vlan="1" port-number="2" enabled="true" layer1-interface-type="unknown" />
                        <switch-port failed="false" managed="false" netboot-enabled="false"
vlan="1" port-number="3" enabled="true" layer1-interface-type="unknown" />
                        <switch-port failed="false" managed="false" netboot-enabled="false"
vlan="1" port-number="4" enabled="true" layer1-interface-type="unknown" />
```

```
                        <switch-port failed="false" managed="false" netboot-enabled="false"
vlan="1" port-number="5" enabled="true" layer1-interface-type="unknown" />
                        <switch-port failed="false" managed="false" netboot-enabled="false"
vlan="1" port-number="6" enabled="true" layer1-interface-type="unknown" />
                        <switch-port failed="false" managed="false" netboot-enabled="false"
vlan="1" port-number="7" enabled="true" layer1-interface-type="unknown" />
                        <switch-port failed="false" managed="false" netboot-enabled="false"
vlan="1" port-number="8" enabled="true" layer1-interface-type="unknown" />
                </switch-module>
                <switch-vlan vlan-number="1" />
        </switch>
        <subnetwork name="ITSOVE" ipaddress="10.1.1.0" netmask="255.255.255.0">
                <vlan vlan-number="1" fabric="ITSO_Fabric" />
        </subnetwork>
        <software-category name="SUSE Linux" />
        <software name="SUSE Linux Server" locale="en_US" type="Operating System"
version="5.2.3790 SP1" package_path="package" category="SUSE Linux" />
        <software-category name="TIO" />
        <software name="TIO" locale="en_US" type="Software" version="2.1"
package_path="package" category="TIO" /> <spare-pool name="ITSOVE Linux" locale="en_US"
os-type=" Linux" vlan="1" fabric="ITSO_Fabric">
                <server name="velin11" locale="en_US">
                        <nic name="Local Area Connection" macaddress="000D609C39D4"
netboot-enabled="false" managed="true" connected-to-switch="ITSO_Switch"
connected-to-port="1">
                                <network-interface name="Local Area Connection"
ipaddress="10.1.1.176" netmask="255.255.255.0" />
                        </nic>
                        <property name="platform" value="IBM eServer HS20"/>
                        <sap name="scp client" port="22" host="false" app-protocol="SCP"
protocol-type="ipv4" locale="en_US">
                                <credentials search-key="primary" is-default="true">
                                        <password-credentials username="root"/>
                                </credentials>
                        </sap>
                        <sap name="ssh service" port="22" host="true" app-protocol="SSH"
is-device-model="SSH Service Access Point" protocol-type="ipv4" locale="en_US">
                                <default-sap operation-type="execute-command"/>
                                <credentials search-key="secondary" is-default="false">
                                        <password-credentials username="tioadmin"
password="password"/>
                                </credentials>
                                <credentials search-key="primary" is-default="true">
                                        <rsa-credentials username="tioadmin"/>
                                </credentials>
                        </sap>
                        <sap name="snmp-set" port="161" host="true" app-protocol="SNMP"
is-device-model="SNMP V1 Service Access Point" protocol-type="ipv4" context="write"
locale="en_US">
                                <default-sap operation-type="set-attribute"/>
                                <credentials search-key="primary" is-default="true">
                                        <snmp-credentials community="public"/>
                                </credentials>
                        </sap>
                        <sap name="snmp-get" port="161" host="true" app-protocol="SNMP"
is-device-model="SNMP V1 Service Access Point" protocol-type="ipv4" context="read"
locale="en_US">
                                <default-sap operation-type="get-attribute"/>
<credentials search-key="primary" is-default="true">
                                        <snmp-credentials community="public"/>
```

```
                                </credentials>
                        </sap>
                        <software-association software="SUSE Linux Server"
state="running"/>
                        <software-association software="TIO" state="running"/>
                </server>
        </spare-pool>
        <kanaha-config>
                <dcm-object id="0">
                        <property component="DEPLOYMENT_ENGINE"
name="PackageIntermediateHost" value="velin11"/>
                        <property component="DEPLOYMENT_ENGINE"
name="PackageRepositoryHost" value="velin11"/>
                        <property component="DEPLOYMENT_ENGINE" name="PackageWebHost"
value="velin11"/>
                        <property component="DEPLOYMENT_ENGINE"
name="SoftwareRepositoryHost" value="velin11"/>
                </dcm-object>
        </kanaha-config>
</datacenter>
```

6

Installing the Virtualization Engine console, IBM Director, TPM, and EWLM on pSeries with AIX

In this chapter, we describe how we installed the following Virtualization Engine components in an AIX environment:

- The IBM Virtualization Engine console
- The IBM Director Server on Windows 2003 and the IBM Director Agent
- The Virtualization Engine console bridge to communicate with the Director Server
- The Enterprise Workload Manager
- Tivoli Provisioning Manager

This chapter discusses the following topics:

- Before the installation
- Installing the Virtualization Engine console component
- Installing the IBM Director Server and Agent components
- Installing the Enterprise Workload Manager component
- Installing Tivoli Provisioning Manager on top of the Virtualization Engine components on a pSeries with AIX

© Copyright IBM Corp. 2005. All rights reserved.

6.1 Before the installation

Before the actual installation, we describe the topology, the installation flow, and the few pre-installation steps that you need to consider.

6.1.1 The topology

In our environment, we used the following configuration, as illustrated in Figure 6-1:

- Management servers
  - The Virtualization Engine console and the Enterprise Workload Manager Management (EWLM) domain manager are installed on the AIX node.
  - The IBM Director Server does not support AIX. We used one Windows 2003 node instead.
- Managed servers

  We have six managed servers with the following operating systems: AIX, Intel Linux, and Windows.
  - EWLM is installed on one AIX managed server and one Windows 2003 managed server. Therefore, if we configure HTTP server or WebSphere Application Server to use ARM, EWLM Domain Manager can monitor those transactions.
  - The IBM Director Agents are installed on all the managed servers. Therefore, you can monitor and manage AIX, Linux, and Windows from one single IBM Director Console.

Note: The name of the box in the figure is the first part of the host name that is used in the network configuration.

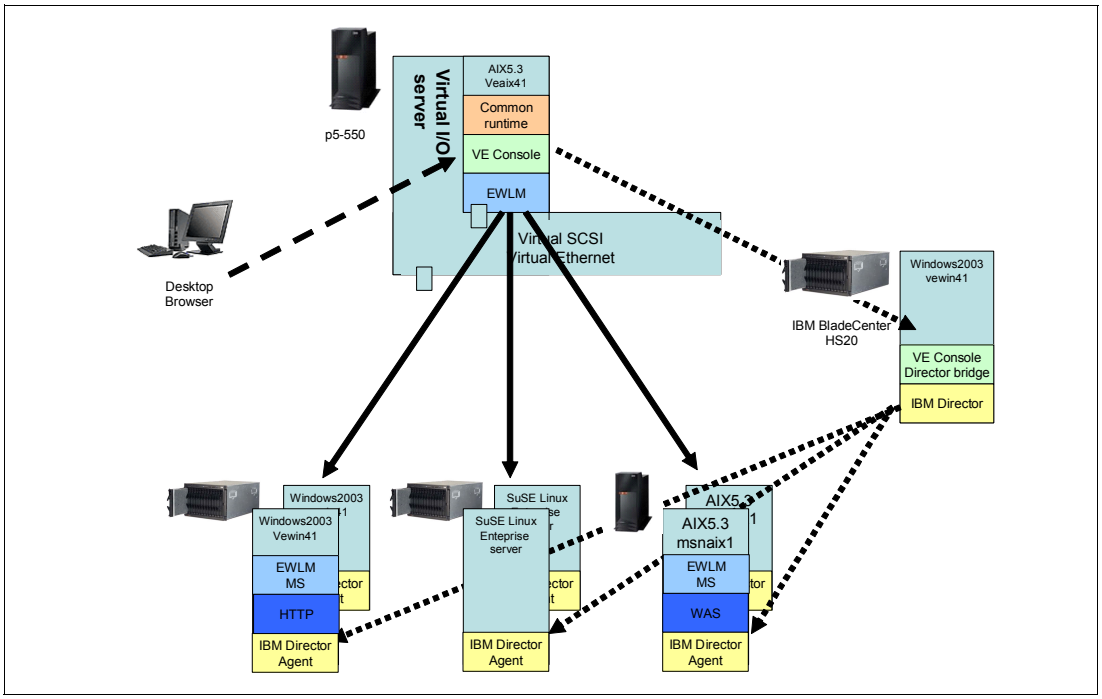

Figure 6-1 The pSeries installation topology

Table 6-1 describes the network configuration and the role of each server in our environment.

Table 6-1 Network configuration and role

Host Name	Operating System	IP Address	Role
veaix41.itso.ibm.com	AIX V5.3	10.1.1.186	EWLM Domain Manager Virtualization Engine console
vewin41.itso.ibm.com	Windows 2003	10.1.1.181	Virutalization Engine Console bridge for IBM Director IBM Director Server
msnaix1	AIX V5.3	10.1.1.189	EWLM Managed Server IBM Director Agent
msnaix2	AIX V5.3	10.1.1.190	IBM Director Agent
Other managed servers	Windows 2003 SLES 8		

6.1.2 The installation flow

The main steps of the installation process in this configuration are the following:

1. Virtualization Engine installation, which includes the following steps:
   a. Installing and configuring prerequisites: the Tivoli Directory Server using the graphical installer, which installs a WebSphere and a DB2 database.
   b. Installing and configuring the Virtualization Engine console code using the Virtualization Engine installer.
2. IBM Director installation of the Server and the Agent components as follows:
   a. Installing a DB2 with the fix pack.
   b. Installing the IBM Director Server code (on Windows) and the IBM Director Agent code (on AIX).
   c. Installing the Virtualization Engine console bridge for the IBM Director to connect to the management source.
   d. Configuring the Virtualization Engine console to display the information that is provided by IBM Director.
3. Enterprise Workload Manager installation, which includes the following steps:
   a. Installing the EWLM Domain Manager using the Virtualization Engine installer.
   b. Configuring the EWLM Domain Manager, using the provided EWLM commands and the configWizard.
   c. Configuring the Control Center according to the required roles we want to install.
   d. Installing and configuring the EWLM Managed Server components, using the configWizard.
4. Tivoli Provisioning Manager installation on top of the previous components.

The following sections describe these steps in more detail.

6.1.3 Pre-installation recommendations

Before you start the installation, we consider the topics presented in this section.

The Virtualization Engine installer

The Virtualization Engine provides a graphical tool installer called Virtualization Engine installer. The Virtualization Engine installer finds the components that are located on disk or CD and shows you the list of them. During the installation process, we choose components that we want to install from the list. We can select all and install them at one time. In the process described here, we installed them separately.

Getting the code

Before starting the installation process, you need to get the installation images. We used files downloaded from a server.

AIX installation methods

On AIX, there are two different methods to install the Virtualization Engine console:

- Using a graphical installer (such as other platforms) is the easiest way.
- Using AIX utilities (such as **geninstall** or SMIT).

In order to use the graphical installer, you must have one Windows server™ on your AIX machine or somewhere on your network. In case of an LPAR environment, it is possible that we do not have a graphical display or Windows. In this case, you can install using **geninstall** or SMIT.

In this chapter, we describe the installation process using the graphical installer. If you are planning to install using AIX utilities, refer to *Using AIX tools to install Virtualization Engine* which is the part of Virtualization Engine chapter of the IBM @server Software Information Center:

```
http://publib.boulder.ibm.com/infocenter/eserver/v1r1/en_US/info/veicinfo/
eicargeninstall.htm
```

In order to use the graphical installer, we need to set the DISPLAY variable in the console which you are going to use for installation. If you log in to the X Window System locally, this variable can be set automatically. If you did not have local graphic display and you are planning to use a remote X Window System server, you need to set this variable by your own. In order to confirm the current setting of the DISPLAY variable, execute the **echo** command. Example 6-1 shows the output.

Example 6-1 confirmation of DISPLAY variable

```
echo $DISPLAY
9.12.6.163
```

If the variable is not set correctly, execute the following command:

```
# export DISPLAY=xserver.itso.ibm.com:0.0
```

In this command, *xserver.itso.ibm.com* is the host name of Windows server. The graphics are displayed on the console of this server. To test your setting, execute the **# xterm** command.

6.2 Installing the Virtualization Engine console component

Before you install the Virtualization Engine console, you must configure an LDAP server. The Virtualization Engine console uses the LDAP server for user and group authentication purpose. The Information Center provides a worksheet of installation and configuration settings to help the installation process.

Tip: Before you begin, we recommend that you complete the worksheet and that you decide the parameters.

Table 6-2 describes the worksheet that we used to install the Virtualization Engine console.

Table 6-2 Virtualization Engine installation worksheet

Virtualization Engine Installation Worksheet	
LDAP server connection information	
Installation screen label	**Answer**
LDAP server host	veaix41.itso.ibm.com
LDAP server port	389 (default)
DN used for directory updates (Administrator DN)	cn=root
Password associated with DN used for directory updates	password
LDAP security settings	
Installation screen label	**Answer**
DN of user and group tree (Suffix DN)	o=ibm, c=us
user naming attribute	uid
Relative location of users	cn=users
User object class	inetOrgPerson
Group naming attribute	cn
Relative location of groups	cn=groups
Group object class	groupOfUniqueNames
Group membership attribute	uniqueMember
User bind information for LDAP server	
Installation screen label	**Answer**
DN for user and group lookup	uid=iscamin, cn=users, o=ibm, c=us
Password associated with DN for user and group lookup	password

Virtualization Engine Installation Worksheet	
LTPA security settings	
Installation screen label	**Answer**
LTPA password	passowrd
LTPA time-out	120
SSO domain name	itso.ibm.com
TCP connection information for the ISC	
Installation screen label	answer
ISC host name	veaix41.itso.ibm.com
Starting port	8421
Integrated Solution Console - LDAP security	
Installation screen label	**Answer**
Portal Administrator DN	uid=iscadmin, cn=users, o=ibm, c=us
short user name	iscadmin
Portal administrator password	password
Portal administrator group DN	cn=iscadmins, cn=group, o=ibm, c=us
Short group name	iscadmins

6.2.1 Installing and configuring Tivoli Directory Server

The Virtualization Engine requires that Tivoli Directory Server V5.1 is installed. The LDAP server can be on a separate machine. If you are planning to use an LDAP server that is already on your network, it must be Tivoli Directory Server V5.1 for AIX environment. Tivoli Directory Server V5.1 is only supported as LDAP server on an AIX environment.

Planning installation tips

Consider the following tips as you plan your installation:

- Ensure that your environment satisfy the system and software requirements of Tivoli Directory Server V5.1

Note: At the time of writing, the Tivoli Directory Server V5.1 does not support AIX V5.3 and the Virtualization Engine console does not support the Tivoli Directory Server V5.2. Therefore, if you are planning to set up AIX V5.3 server as a management server, you need another server for Tivoli Directory Server V5.1.

- Ensure that the /tmp file system has enough free space to expand DB2 image.

  If you do not have enough space, the installation will fail. In our environment, it is 300 MB.

- Check the status of an asynchronous I/O device.

  An asynchronous I/O device must be available. If its status is defined, execute the **`smit aio`** command and select Configure Defined Asynchronous I/O.

In order to make aio available automatically when the system is restarted, execute the **smit chgaio** command and change the STATE to be configured at system restart parameter to available, as shown in Figure 6-2.

```
Change / Show Characteristics of Asynchronous I/O

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                                  [Entry Fields]
  MINIMUM number of servers                      [1]                          #
  MAXIMUM number of servers per cpu              [10]                         #
  Maximum number of REQUESTS                     [4096]                       #
  Server PRIORITY                                [39]                         #
  STATE to be configured at system restart        available                   +
  State of fast path                              enable                      +
```

Figure 6-2 SMIT example of an I/O setting

- The Tivoli Directory Server Rebidding Tool conflicts with the AIX Web-based System Manager. Both utilities use the same TCP/IP port (9090). If the AIX Web-based System Manager is started, stop it during installation and configuration. To stop the Web-based System Manager, execute the following command:

  ```
  # /usr/websm/bin/wsmserver -disable
  ```

  If you need to enable Web-based System Manager after LDAP installation, execute the following command:

  ```
  # /usr/websm/bin/wsmserver -enable
  ```

- DB2 UDB is installed with the Tivoli Directory Server. We need to create a DB2 database owner before starting the installation. Example 6-2 shows the procedure we used to create a DB2 user.

Example 6-2 Example of user creation for DB2

```
# mkuser db2admin
# passwd db2admin
Changing password for "db2admin"
db2admin's New password:
Enter the new password again:
# pwdadm -f NOCHECK db2admin
# mkgroup -a users=root,db2admin dbsysadm
# chuser pgrp=dbsysadm db2admin
# chown db2admin.db2sysadm /home/db2admin
# chown db2admin.dbsysadm /home/db2admin
```

Installing the Tivoli Directory Server V5.1 code

The graphical installer is provided to install Tivoli Directory Server V5.1 on AIX. If your environment does not have a GUI environment, you can use SMIT or the **installp** command to install it. The installation process with the graphical installer is well described in the Information Center, but the SMIT and **installp** installation processes are not described. If you are planning to install with SMIT or **installp**, refer to the Tivoli Directory Server V5.1 Installation and Configuration Guide.

In our installation, we used the graphical installer and followed the steps described in the IBM @server Software Information Center.

The installer installs WebSphere Application Server - Express V5.0 and DB2 Universal Database V8.1. WebSphere Application Server is installed in the /usr/ldap/appsrv directory. DB2 Universal Database is installed in /usr/opt/db2_08_01. The LDAP server components are installed in the /usr/ldap directory. These directories cannot be changed during the installation.

Tip: The installation log file is named /usr/ldap/ldapinstall.log. If you have any problem during the installation, you can find detailed failure reasons in this file.

When you finish the installation, the Tivoli Directory Server Configuration Tool runs automatically (see Figure 6-3). If you want to configure later, you can close this window. You can start it by using the following command:

```
# /usr/ldap/bin/ldapxcfg
```

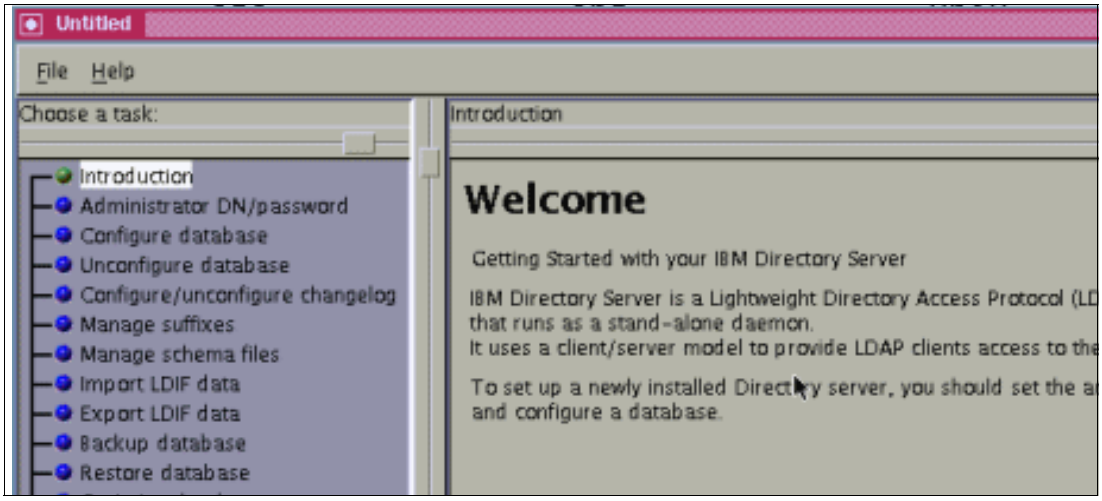

Figure 6-3 The welcome page of ldapxcfg

We set up the LDAP Administrator DN and the password, and we configured the database using this tool as follows:

- In order to set the Administrator DN and the password, click **Administrator DN/password**. The LDAP server connection information section on Table 6-2 on page 121 provides this information. In our environment, the administrator DN is `cn=root`, and the password is `password`.
- In order to create a database, click **Configure database**. When we create the database, we need to provide a DB2 database owner user and the password. In our environment, the DB2 database owner is `db2admin`, and the password is `password`. We followed the instruction of the Information Center for the other settings.

A command line interface **`ldapcfg`** is also provided. In order to set the Administrator DN and password using this interface, execute the following command:

```
# ldapcfg -u"cn=root" -p"password"
```

To create the database, execute the following command:

```
# ldapcfg -c -a db2admin-w password -d ldapdb -l /home/db2admin
```

In these commands:

- **-a** is the database administrator user name
- **-w** is the database administrator user password
- **-d** is the database name
- **-l** is the database directory
- **-c** create a database in UTF-8

Configuring the LDAP directory for the Virtualization Engine console

After we have installed the LDAP server and we have finished the basic configuration of the LDAP server, we are ready to begin setting up our LDAP server for the IBM Virtualization Engine console. For the LDAP configuration, the Tivoli Directory Server Web Administration Tool is provided with the Tivoli Directory Server. This is a Web application running on the WebSphere Application Server - Express. The detailed information is described in the Information Center. In this section, we followed the instruction from the Information Center.

To configure the LDAP directory:

1. Start the Tivoli Directory Server using the **ibmdirctl** command, as shown in Example 6-3.

Example 6-3 Starting LDAP server with ibmdirctl command

```
# ./ibmdirctl -D"cn=root" -w? start
Enter password ==>
Start operation succeeded
#
```

2. Start the Web Administration Tool. The Web Administration Tool is a Web application running on WebSphere Application - Express. Therefore, in order to start Web Administration Tool, you need to start WebSphere Application Server. WebSphere Application Server - Express is installed in the /usr/ldap/apprv directory. Example 6-4 provides the command to start the Web Administration Tool.

Example 6-4 Starting the Web Administration Tool

```
# cd /usr/ldap/appsrv/bin
# ./startServer.sh server1
ADMU0116I: Tool information is being logged in file
           /usr/ldap/appsrv/logs/server1/startServer.log
ADMU3100I: Reading configuration for server: server1
ADMU3200I: Server launched. Waiting for initialization status.
ADMU3000I: Server server1 open for e-business; process id is 45536
#
```

Tip: Ensure that AIX Web-based System Management is stopped. If it is not stopped, WebSphere Application Server will fail to start because of duplication of port requests.

3. Open your browser, and go to the following URL, where *hostname* is the name of your LDAP server:

   `http://hostname:9080/IDSWebApp/IDSjsp/Login.jsp`

   Figure 6-4 shows a screen shot of the login page. In order to start the configuration, you need to add your server on this console first. Select **Console Admin as LDAP Hostname**, and log in as a user `superadmin` using the password `secret`. This is the default user for the console.

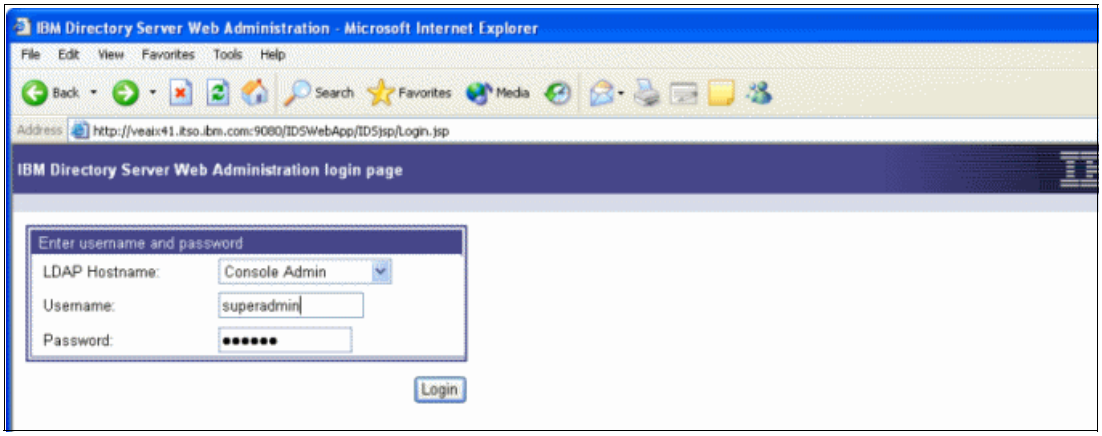
IBM Directory Server Web Administration - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address http://veaix41.itso.ibm.com:9080/IDSWebApp/IDSjsp/Login.jsp
IBM Directory Server Web Administration login page
Enter username and password
LDAP Hostname: Console Admin
Username: superadmin
Password:
Login

Figure 6-4 Login page of Tivoli Directory Server Web Admin Tool

4. When you log in, expand Console Administration and click **Manage console servers**. Then, click **Add** and enter your server name, port, and SSL configuration as described in the Information Center.
5. Configure the LDAP directory by using the Web Administration Console. (We are in N4 of step 4 in the instructions as provided by the Information Center.) Log off from the Console Admin, and log in to your server console as the LDAP server administrator. In our environment, the LDAP host name is named `veaix41.itso.ibm.com`, the username is `cn=root`, and the password is `password`.
6. Create the suffix DN entry of the LDAP tree for the ISC user and group. In our environment, the suffix DN is `o=ibm, c=us`. Thus, we select organization as a structual object class.
7. Create container nodes which store user and group objects. These nodes are Relative Location of groups and Relative Location of users value. In our environment, we create `cn=users` and `cn=groups`.
8. Add users to the LDAP directory. We add administrator user `uid=iscadmin, cn=users. o=ibm, c=us`. Add the administrator group `cn=iscadmins, cn=group, o=ibm, c=us`.

The installation of the LDAP is now finished.

Tivoli Directory Server does not starts automatically at boot time by default. If you want it to start automatically, use the following command.

```
# /usr/sbin/mkitab "ldap:2:once:/usr/ldap/bin/slapd -f /etc/ibmslapd.conf \
> /dev/console 2>&1"
```

This command edits the /etc/inittab file.

Verify the installation

In order to verify the LDAP server installation, use the **ldapsearch** command. Example 6-5 shows the command and the output. This command shows the attributes that are filtered by uid=iscadmin. If you enter cn=* instead of uid=iscadmin, you can see all the objects that have cn.

Example 6-5 An ldapsearch example

```
# ldapsearch -b"o=ibm, c=us" "uid=iscadmin"
uid=iscadmin,cn=users,o=ibm,c=us
objectclass=inetOrgPerson
objectclass=organizationalPerson
objectclass=person
objectclass=top
objectclass=ibm-appuuidaux
sn=iscadmin
cn=iscadmin
uid=iscadmin
ibm-appuuid=a6e77bb0-c0cc-11d9-9dcd-8282427fce37
```

6.2.2 Installing the Virtualization Engine console code

After we installed and configure the LDAP server, we install the Virtualization Engine console code. The installation process for the Virtualization Engine console is simple if you complete the worksheet. The Virtualization Engine installer guides us.

Planning installation tips

Before you begin to install the Virtualization Engine console ensure the following:

- That your environment satisfies system requirements. They are described and updated in the Information Center.
- That the LDAP server has been installed and configured correctly.
- That the LDAP server is started.
- That the installation target directory and the /tmp directory have enough space to install the Virtualization Engine console. The installer does not expand these directories automatically. The installer tells you how much free space is required if there is not enough space.

Installing the code

We can use the Virtualization Engine installer to install the Virtualization Engine console. In order to start the installer, enter the following commands;

```
# cd <your install image directory>
# ./installVEAix.sh
```

After you start the installer, execute the following steps;

1. Accept the license agreement. The installer asks you where you install software. The default directory is /opt/IBM/VE. Select the components that you are going to install. We installed only the Virtualization Engine console (Figure 6-5).

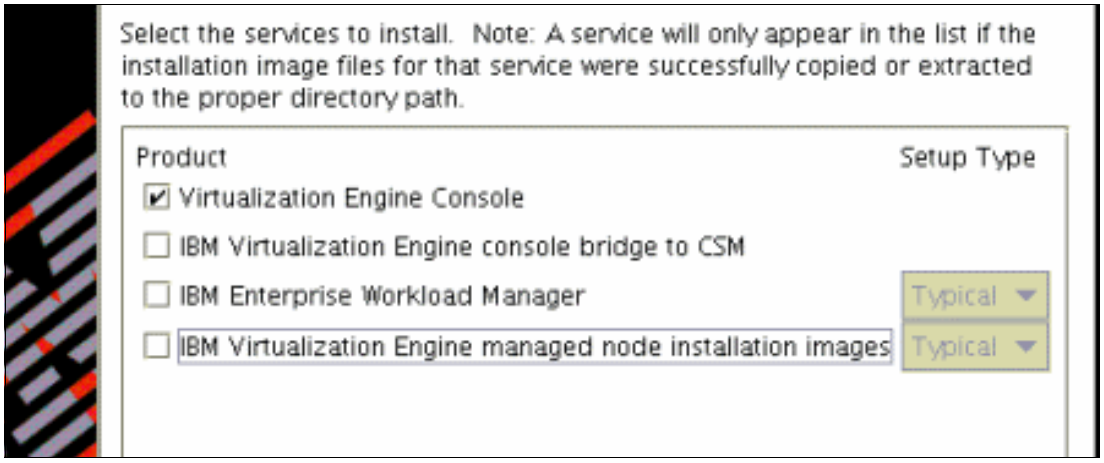
Select the services to install. Note: A service will only appear in the list if the installation image files for that service were successfully copied or extracted to the proper directory path.
Product
Setup Type
Virtualization Engine Console
IBM Virtualization Engine console bridge to CSM
IBM Enterprise Workload Manager
Typical
IBM Virtualization Engine managed node installation images
Typical

Figure 6-5 Package selection

2. The installer next asks about the LDAP settings. You need to complete the LDAP connection information. Each section on Table 6-2 on page 121 corresponds to the installation screen.

   For example, according to the LDAP connection information, the LDAP server host is `veaix41.itso.ibm.com`, the LDAP server port is `389`, the domain name that is used for directory updates is `cn=root`, and the password is `password` as shown in Figure 6-6.

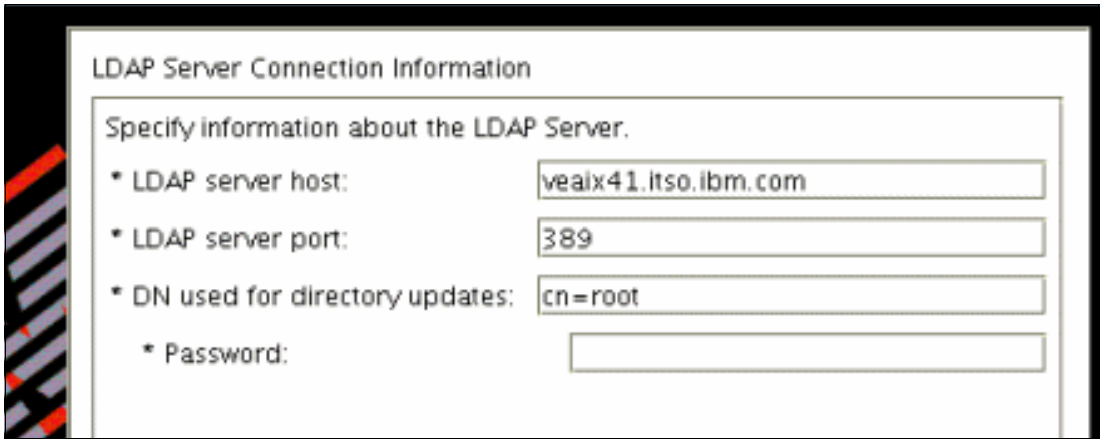

Figure 6-6 LDAP Server Connection Information

3. Complete the LDAP security settings and the user bind information as shown in Figure 6-7 and Figure 6-8.

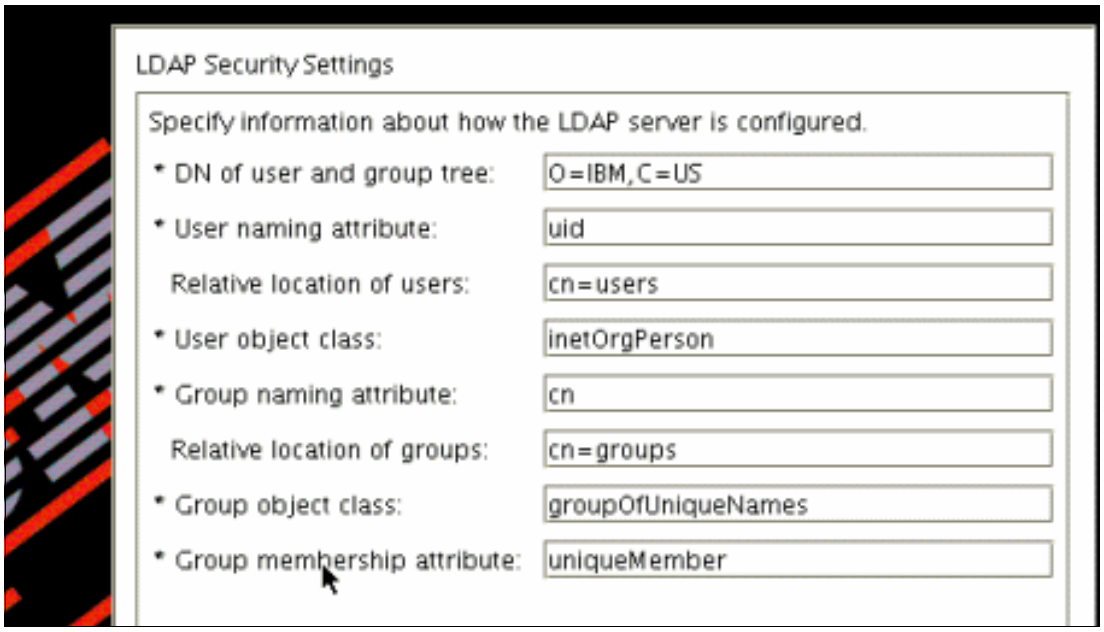

Figure 6-7 LDAP Security Settings

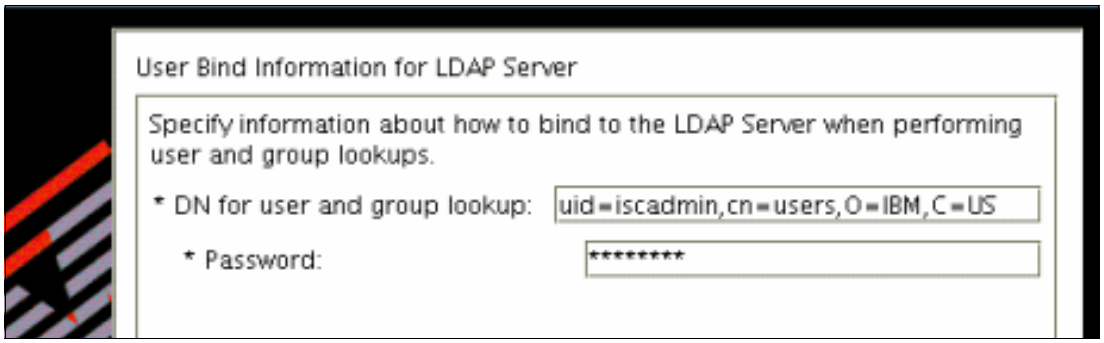
User Bind Information for LDAP Server
Specify information about how to bind to the LDAP Server when performing user and group lookups.
* DN for user and group lookup:
uid=iscadmin,cn=users,O=IBM,C=US
* Password:
****

Figure 6-8 User Bind Information for LDAP Server

4. The installer asks for the LTPA key information, as shown in Figure 6-9. Complete the LTPA password, timeout, and SSO domain information.The LTPA key file and password are required to install the Virtualization Engine bridge for each management source (for example, IBM Director and CSM). In this step, you create and export the LTPA key file. We need this file to install the Virtualization Engine console bridge for IBM Director as shown in Figure 6-10.

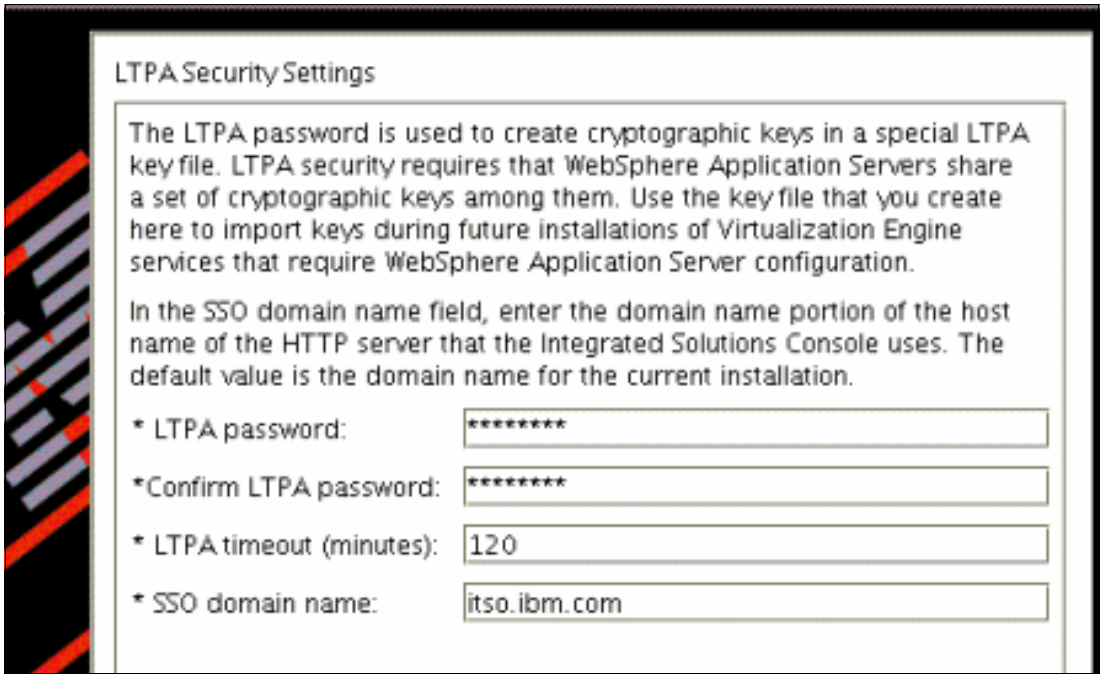
LTPA Security Settings
The LTPA password is used to create cryptographic keys in a special LTPA key file. LTPA security requires that WebSphere Application Servers share a set of cryptographic keys among them. Use the key file that you create here to import keys during future installations of Virtualization Engine services that require WebSphere Application Server configuration.
In the SSO domain name field, enter the domain name portion of the host name of the HTTP server that the Integrated Solutions Console uses. The default value is the domain name for the current installation.
* LTPA password:
****
*Confirm LTPA password:
****
* LTPA timeout (minutes):
120
* SSO domain name:
itso.ibm.com

Figure 6-9 LTPA Security Settings

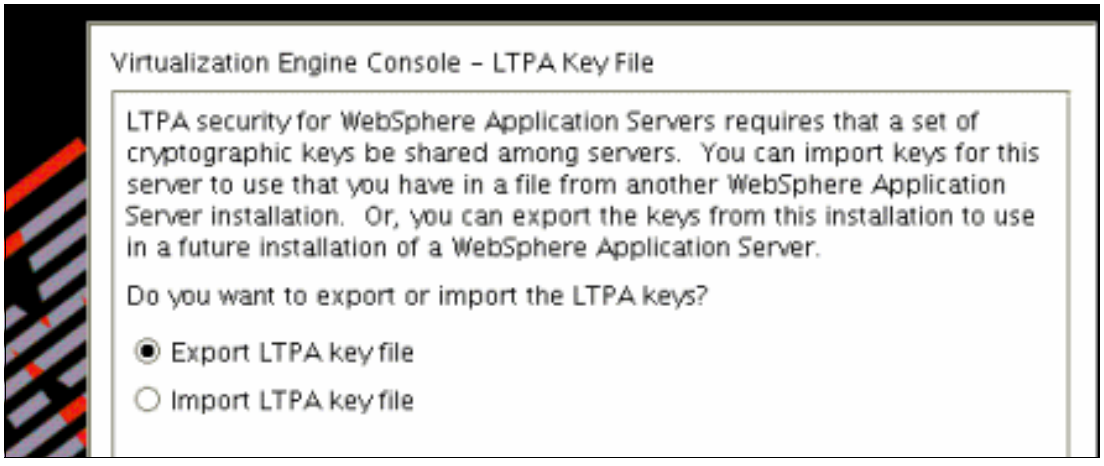
Virtualization Engine Console – LTPA Key File
LTPA security for WebSphere Application Servers requires that a set of cryptographic keys be shared among servers. You can import keys for this server to use that you have in a file from another WebSphere Application Server installation. Or, you can export the keys from this installation to use in a future installation of a WebSphere Application Server.
Do you want to export or import the LTPA keys?
Export LTPA key file
Import LTPA key file

Figure 6-10 Export LTPA key file

5. The installer then asks for the ISC configuration. Each parameter is already completed with the values that you specified. You just need to confirm each parameters. After you enter the starting port for ISC, the installer shows you the Virtualization Engine console ports as shown in Figure 6-11. You need to note these ports.

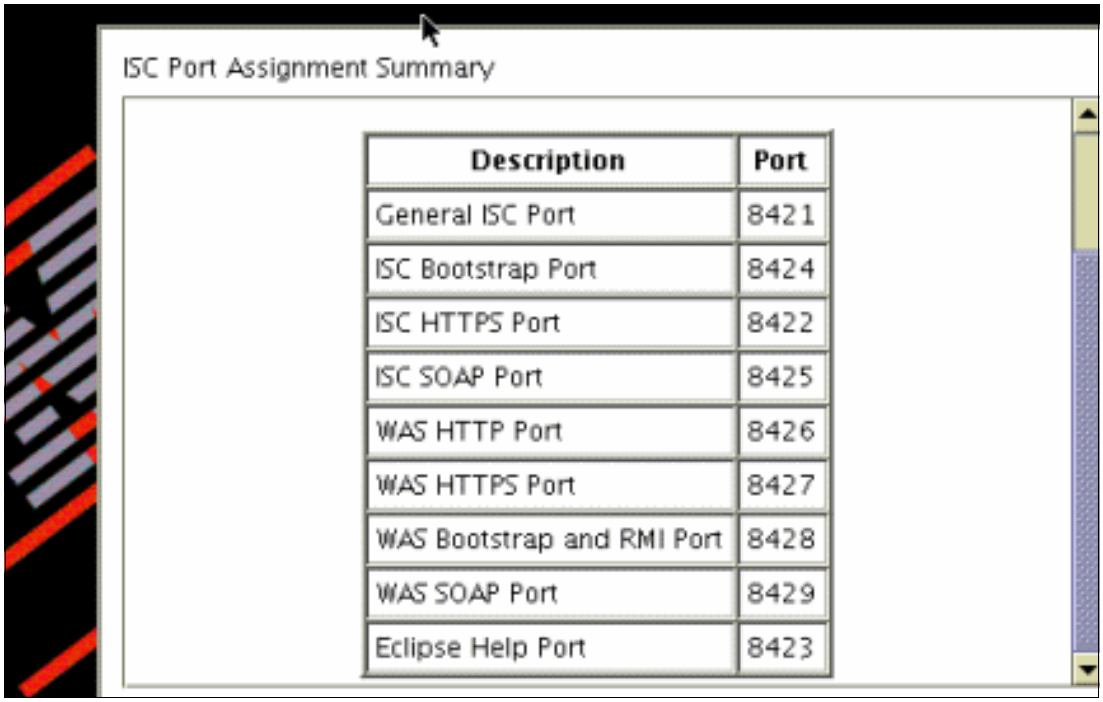

Figure 6-11 Port information

6. The installer shows an installation configuration summary. If you agree with all the parameters, click **Next** to begin the installation. After the installation is complete, the installer asks you whether you want to open the Virtualization Engine console. If you installed the browser in your management node, you can start your browser.

Verifying the installation

In order to verify your installation, access the Virtualization Engine console by using the browser. Follow these steps:

1. Start the Virtualization Engine console by starting the WebSphere Portal Server on which the Virtualization Engine console is defined.

   Log in as root, and then run the following commands:

   ```
   # cd /opt/IBM/VE/ISC/PortalServer/bin
   # ./startISC.sh ISC_Portal
   ```

   When you run this command, you enter `iscadmin` as the administrator user ID for portal and `password` as the password for the administrator user.

2. Access the Virtualization Engine console through the browser. In our environment, the URL for the Virutalization Engine console is:

   ```
   http://veaix41.itso.ibm.com:8421/ibm/console
   ```

Figure 6-12 shows the login page of the Virtualization Engine console.

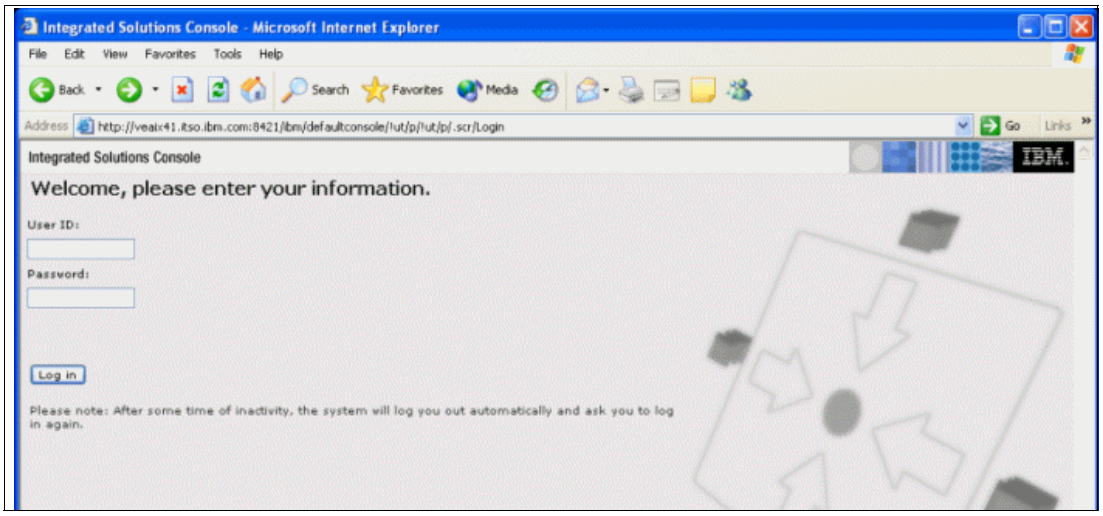
Integrated Solutions Console - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address http://veaix41.itso.ibm.com:8421/ibm/defaultconsole/!ut/p/!ut/p/.scr/Login
Go Links
Integrated Solutions Console
IBM.
Welcome, please enter your information.
User ID:
Password:
Log in
Please note: After some time of inactivity, the system will log you out automatically and ask you to log in again.

Figure 6-12 Login page of Virtualization Engine console

Enter your administrator user ID and password. In our environment, the administrator user ID is `iscadmin` and the password is `password`. If you can log in to the console, the installation has completed successfully.

6.3 Installing the IBM Director Server and Agent components

At the time of writing, the IBM Director Server is not supported on AIX V5.3. Therefore, In order to configure the IBM Director system in this environment, you need at least one Windows 2000/2003 or xLinux machine to install it. In our environment, we used a Windows 2003 operating system to support the IBM Director Server.

The IBM Director Server requires an SQL database system to store the system inventory data. The following SQL database are supported on a Windows platform:

- IBM DB2 Universal Database 8.1, Fix Pack 5
- Microsoft Data Engine (MSDE) 2000, with Service Pack 3a
- Microsoft Jet 4.0 database engine, with Service Pack 8
- Microsoft SQL Server 2000 Desktop Engine, with Service Pack 3a
- Microsoft SQL Server 2000, with Service Pack 3a
- Oracle Server, versions 8.1.7, 9.0, and 9.2

On a Windows platform, we can use a built in database such as Microsoft Jet 4.0 database engine. It is possible to use this database. However, if your environment is a large environment, it is recommended to use another database such as DB2 Universal Database. In our ITSO environment, we used DB2 Universal Database Workgroup Server Edition V8.1.

Before beginning the installation:

- Ensure that your environment satisfies system and software requirements.
- Ensure that a supported database server is installed and configured.
- Build the administration user account before starting the installation. This account must have the administrator privileges on the management server. The IBM Director Server service runs as this account. This user should satisfy the following requirements:
  - Be a member of the administrators group
  - **Password never expires** is selected

6.3.1 Installing DB2 Universal Database V8.1 with Fix Pack 5

In our environment, we use DB2 Universal Database V8.1 as the SQL database for IBM Director. There is no special consideration for IBM Director on installation of DB2. We used db2admin as db2 user, and we followed the installation process as described in the Information Center.

After you install DB2 Universal Database V8.1 base, you need to apply Fix Pack 5 to support the IBM Director. Download the appropriate fix pack from the DB2 support Web site and apply Fix Pack 5 from the following:

```
http://www-306.ibm.com/software/data/db2/udb/support/
```

After the installation is complete, execute the **db2sampl** command to create a sample database. If you can create a sample database, the installation has completed successfully.

You next need to create a database for the IBM Director. It is possible to create a database after installing the IBM Director by using the **cfgdb** command. However, in our environment, we created a database in advance and granted DBADM permission to IBM Director Server. If you select to setup your environment this way, do the following:

1. Log in as a database administrator.
2. Start the database manager.
3. Create a database for IBM Director, using the following command from the DB2 command line processor:

   ```
   create database DIRDB
   ```

4. Grant DBADM authority to the administrator user of IBM Director using the following commands:

   ```
   connect to DIRDB
   grant DBADM on database to user director
   db2 connect reset
   ```

6.3.2 Installing the IBM Director Server code on Windows 2003

Next, you install the IBM Director code. The installation process for IBM Director on this platform is described in 7.3, “Installing the IBM Director components” on page 183.

Verifying the installation

Log in to the console and find the green circle on the right side of the task bar. Then, start the IBM Director Console by selecting **Start** → **All Programs** → I**BM Director** → IBM **Director Console**. The login window is displayed. Enter your director administrator user name and password, and click **OK**. If you can log in successfully, the director console opens, as shown in Figure 6-13 on page 134.

Tip: If you are logging in using Remote Desktop, you might not find a green circle on your task bar. In order to check whether the server is running, open the Control Panel and select **Administrative Tools** → **Services**. Check the status of IBM Director Agent, Server, and Support services.

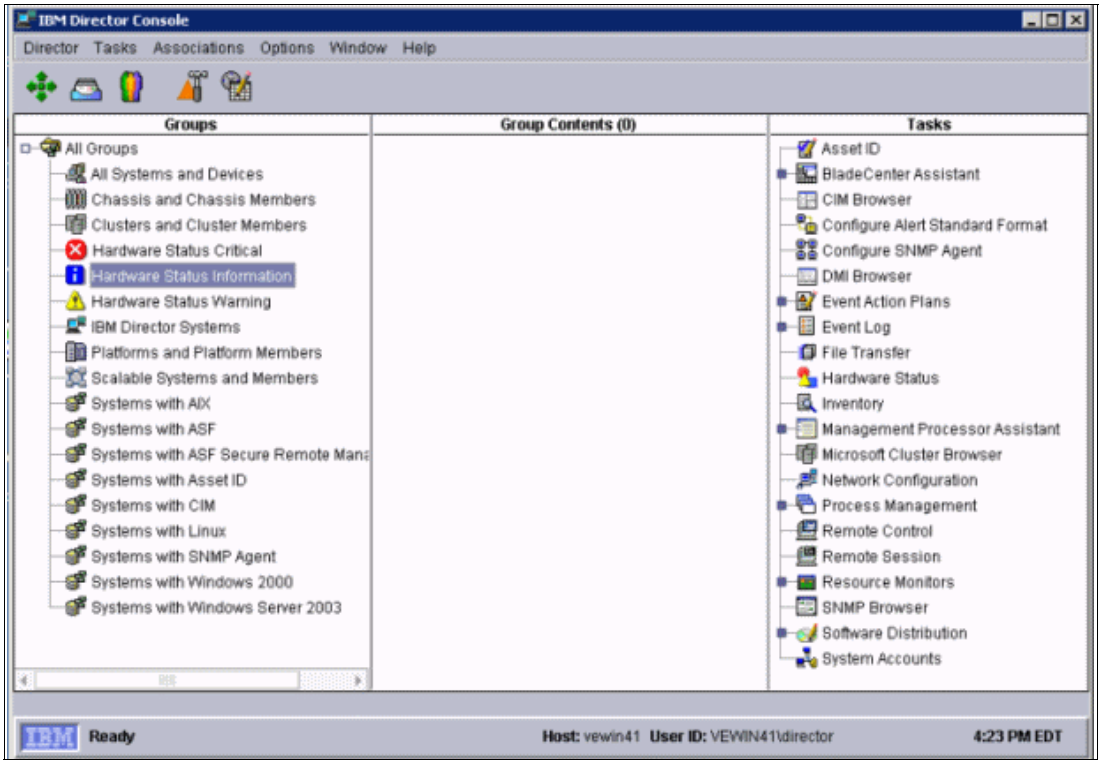

Figure 6-13 IBM Director Console

6.3.3 Installing the IBM Director Agent code on AIX

If you select the *Virtualization Engine managed server installation images* option on the install step, EWLM Managed Server package and IBM Director Agents are installed in *install_directory*/ManagedNode directory for AIX and *install_directory*\ManagedNode for Windows. The package for AIX is the dir4.20_agent_aix.tar file. In our environment, we obtain this package for Windows node and extract dir4.20_agent_aix.tar on our managed server. We used the FTP put method to transfer the file from Windows to AIX.

Launch the IBM Director Agent using the following command:

```
IBMDirectorAgent4.20.bin
```

This is a self-extracting file which runs the **installp** command automatically.

```
# tar -xvf dir4.20_agent_aix.tar
# ./IBMDirectorAgent4.20.bin
```

The IBM Director Agent is installed in the /opt/ibm/director directory. The installation process adds the entry in the /etc/inittab file and starts the agent automatically.

```
IBMDirA:2:wait:/opt/ibm/director/bin/twgstart
```

Verifying the installation

To verify the installation, execute the **ps** command and the **grep** command with **TWGAgent** as a parameter. There must be two JAVA processes for the agent. Then, log in to the IBM Director Console and confirm the connection between the IBM Director Server and the Agent on the IBM Director console. Login to the IBM Director Server or the IBM Director Console node. Start the IBM Director Console. When you start the console, you might not find any system on your console, as shown in Figure 6-14.

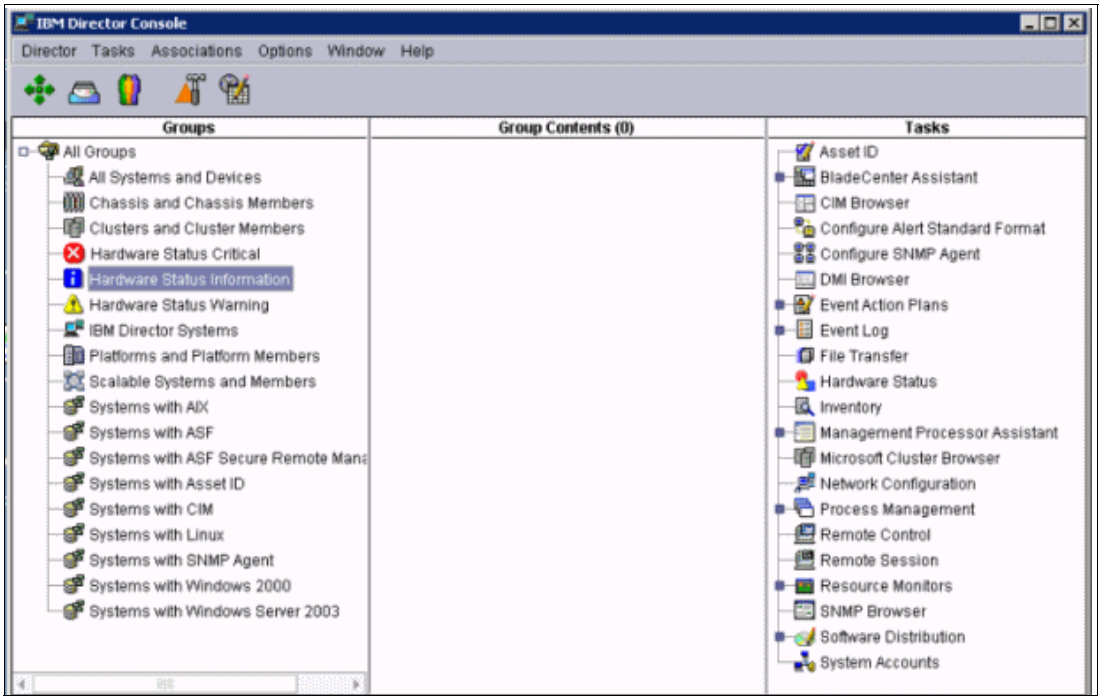

Figure 6-14 IBM Director Console

Right-click **IBM Director Systems** on left panel of the console. Select **Discovery System** → **IBM Director System**. After few seconds, you find your system in the center panel. See Figure 6-15 on page 136 and Figure 6-16 on page 136.

Tip: If you have two or more network interfaces on your server , make sure you know which device is used by the server.

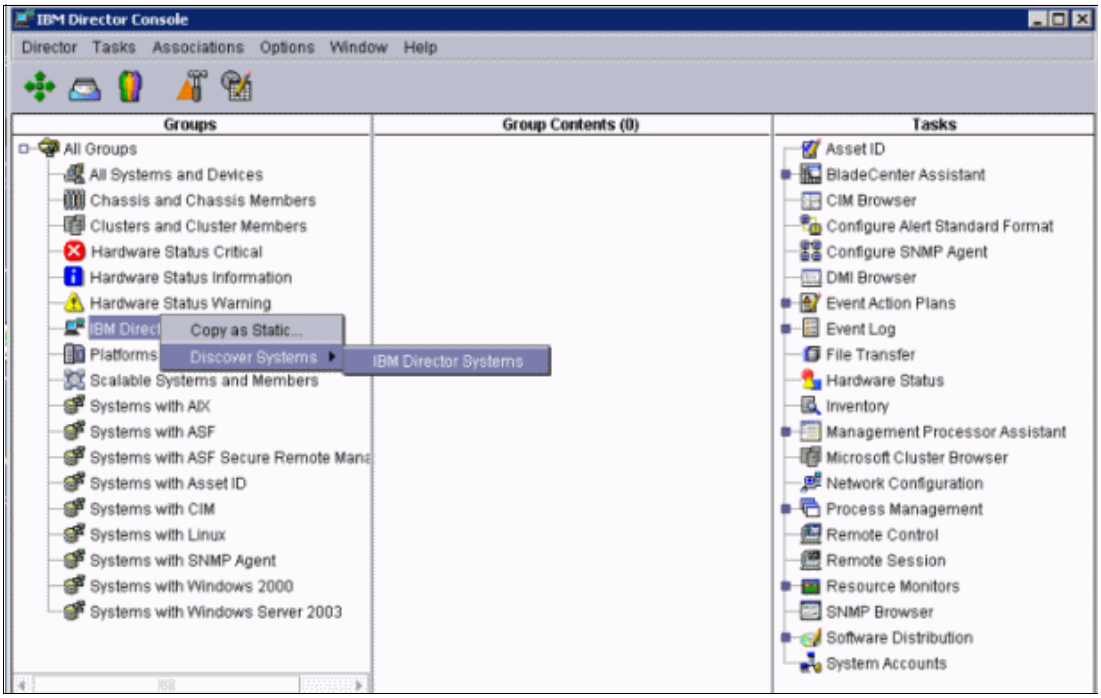

Figure 6-15 IBM Director Console - Device Discovery

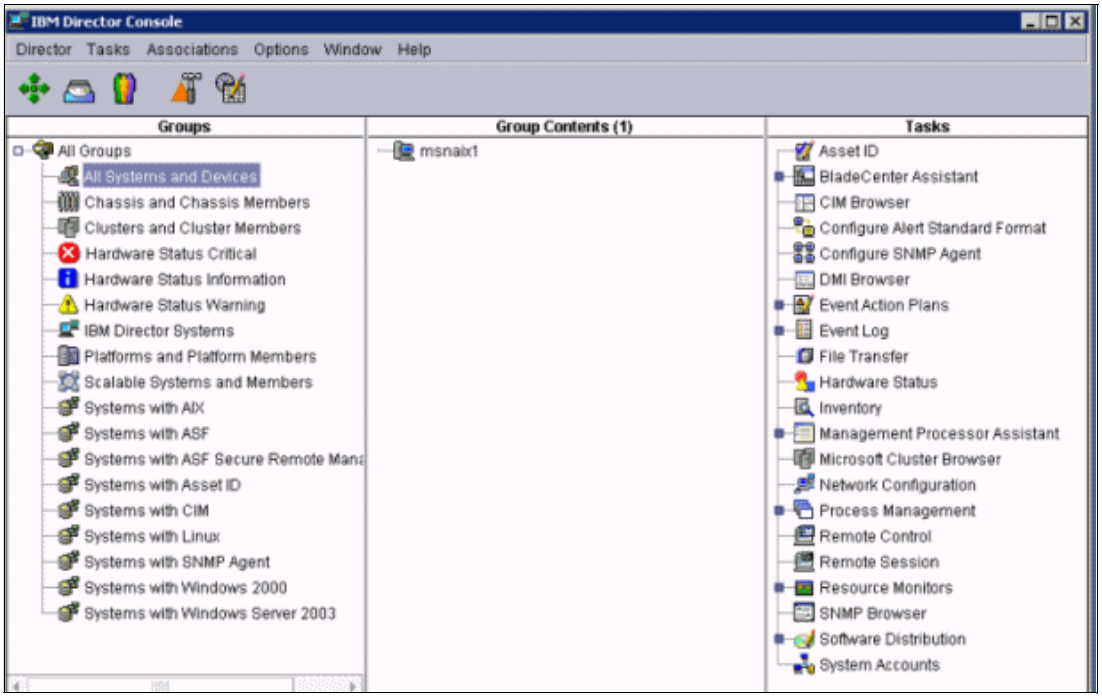

Figure 6-16 IBM Director Console - Discovered Device

You can configure the Virtualization Engine console to communicate with the IBM Director Console.

6.3.4 Installing the Virtualization Engine console bridge for IBM Director

At this point, The Virtualization Engine console and the IBM Director system are already installed and configured. For the Virtualization Engine console to connect to the management source, you must install the Virtualization Engine console bridge for IBM Director. The bridge component is provided for each management source (CSM, IBM Director, and Management Central). The Virtualization Engine console bridge to the IBM Director can be installed on a different physical system from the IBM Director Server. However, in our environment, we install them on the same physical system.

To install this component:

1. Before you begin the installation, prepare the LTPA key which is exported from the Virtualization Engine console installation. We need to share same LTPA key file and password between the Virtualization Engine console and the bridges.
2. To install the Virtualization Engine console bridge for IBM Director, use the Virtualization Engine installer. The installation process is similar to the installation of the Virutalization Engine console.
   a. Select the products to install. We select the Virtualization Engine console bridge for IBM Director and click **Next**.
   b. The installer asks us about LDAP settings. Enter the same value as the Virtualization Engine console installation. Refer to Table 6-2 on page 121 about each value.
3. If the installer requires the LTPA Key, enter the location of the key file. This location must be the file that is exported during the Virtualization Engine console installation.
4. If the installer asks you the installation directory, enter the parent directory for the Virtualization Engine directory. We take the default directory (C:\VECdir) here.
5. If the installer requires the port for WebSphere Application Server that is used by the Virtualization Engine console bridge for IBM Director, enter the appropriate port number. We accept the default port number (12000) here.
6. The installer requires IBM Director Server information. Enter your director server host name and administrator user name and password. If you want to confirm the connection to your IBM Director Server, select **Connect to the IBM Director Server to validate information**.
7. At the end, the installer shows the summary of the installation setting. If it is correct, click **Next** to begin the installation.

Configuring the Virtualization Engine console

To configure the Virtualization Engine console to display the information that is provided by the IBM Director, do the following:

1. Open your browser and go to the following URL to access the Virtualization Engine console:

   `http://hostname:port/ibm/console`

2. Log in as administrator user (for example *iscadmin*).
3. Register the IBM Director Server on the Virtualization Engine console by doing the following:
   a. Select **Virutalization Engine console** → **Preference** → **Management Source**, as shown in Figure 6-17 on page 138.
   b. Click **Add**. The Management Source configuration window opens.
   c. Choose your Management Source type and enter the host name and port number. In our case, we select IBM Director Management source for management source type. If

we select management source type, the default value of port and ContextRoot is entered. Figure 6-18 shows our environment.

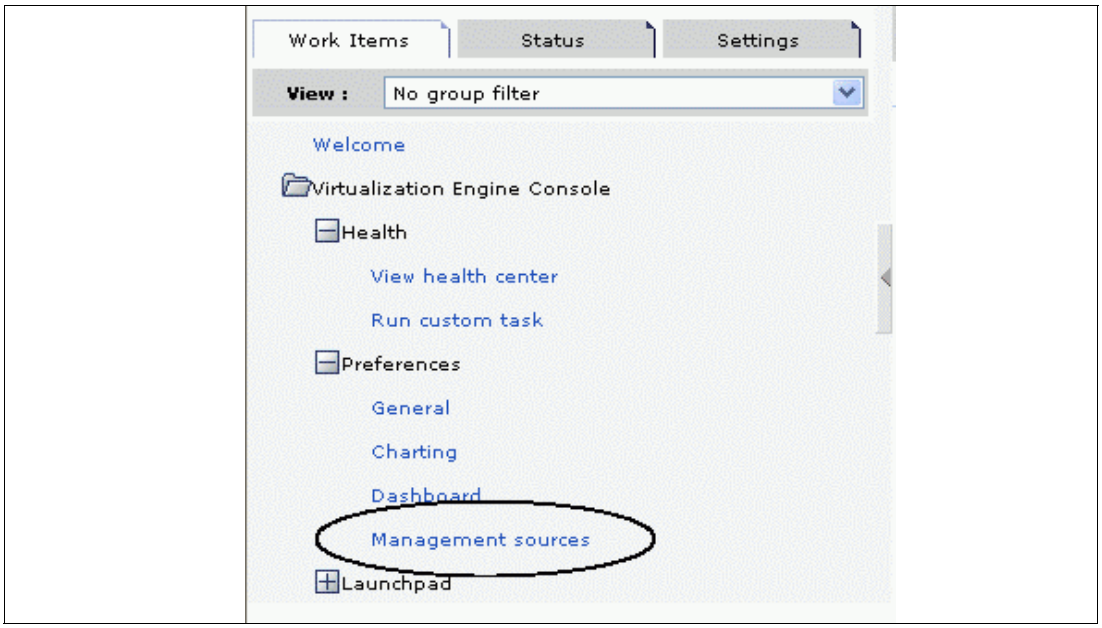

Figure 6-17 Open Management Source

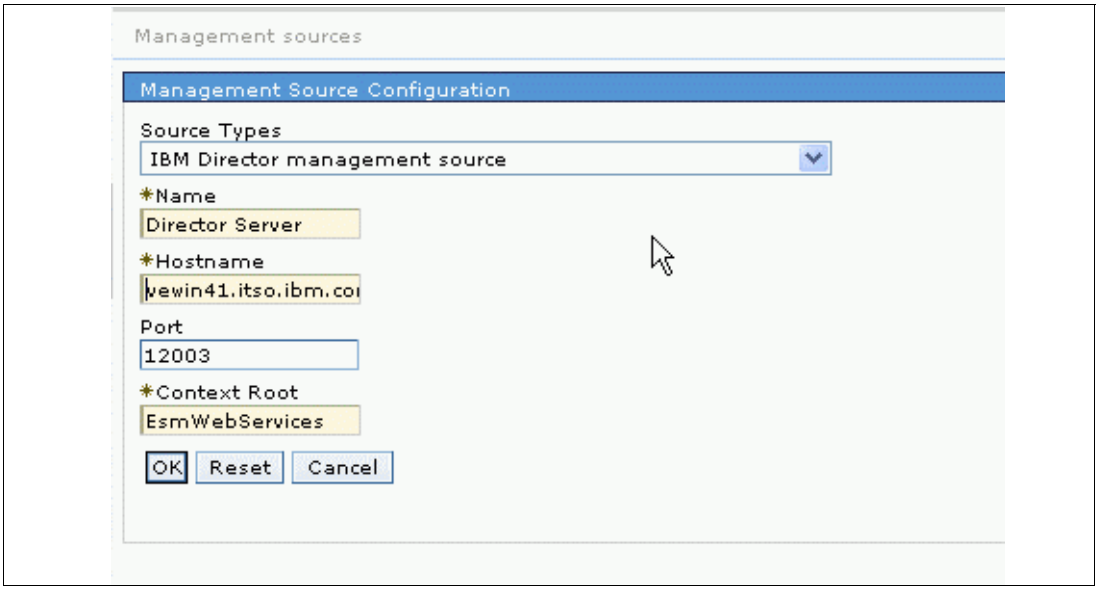

Figure 6-18 Management source configuration

4. Add resources and monitors in the Health Center of the Virtualization Engine console and execute custom tasks. Every resource and monitor must be configured on the management source before you add them. For more information about resources and monitors, refer to 3.4, “IBM Virtualization Engine console” on page 46.

6.4 Installing the Enterprise Workload Manager component

In this section, we describe how to install and configure Enterprise Workload Manager (EWLM) on an AIX system.

6.4.1 Planning the installation

Before you begin the installation process, consider the following:

- Ensure that your environment meets the system requirements. They are described with the latest updates on the Information Center.
- The IBM Software Information Center provides a worksheet for EWLM. Complete this worksheet. For planning information about EWLM, refer to *IBM Enterprise Workload Manager*, SG24-6350.
- Create a WebSphere Admin user.

  During the configuration of EWLM Domain Manager, you need to specify the WebSphere admin user. Example 6-6 shows an example of how to create a user on AIX.

Example 6-6 EWLM admin user creation example

```
# mkuser ewlmadm
# passwd ewlmadm
Changing password for "ewlmadm"
ewlmadm's New password:
Enter the new password again:
# chuser pgrp=system ewlmadm
```

- Decide the location of the working directory for EWLM.

6.4.2 Installing the EWLM Domain Manager

To install EWLM, use the Virtualization Engine Installer. To install the EWLM Domain Manager do the following:

1. You can start the installer by executing the **installVEAix.sh** command similar to the other components.

   We select EWLM and Virtualization Engine Managed Server installation images as install products.

   The installation of Virtualization Engine managed node installation images copies the EWLM Managed Server installation images in the *Install Directory*/ManagedNode/ directory. The typical installation copies the managed server images for all platforms. If you want to copy for specific platforms only, select the custom installation.

   If you want to install single-system EWLM, select a custom installation. The typical installation installs the domain manager only.

   In this section, we select a typical installation for EWLM and a custom installation for Virutalization Engine managed server installation as shown in Figure 6-19 on page 140.

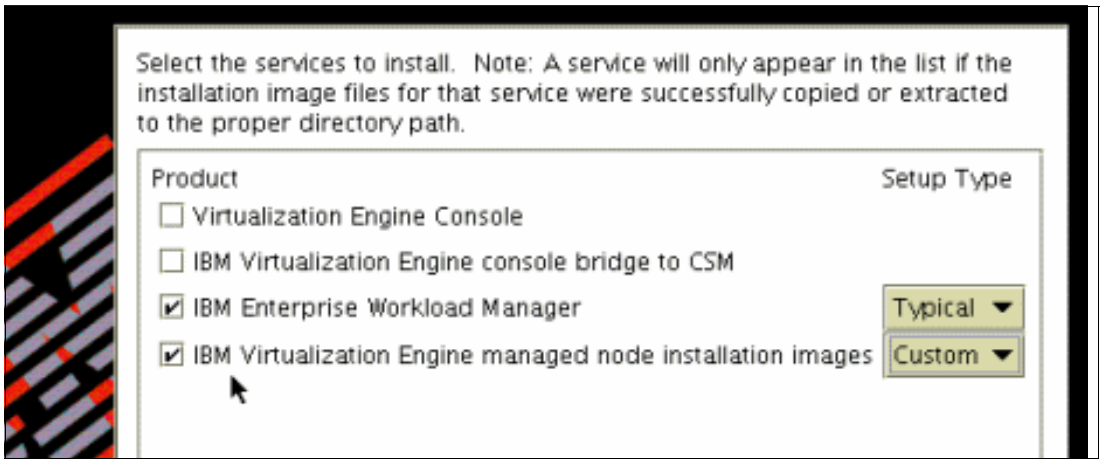
Select the services to install. Note: A service will only appear in the list if the installation image files for that service were successfully copied or extracted to the proper directory path.
Product
Setup Type
Virtualization Engine Console
IBM Virtualization Engine console bridge to CSM
IBM Enterprise Workload Manager
Typical
IBM Virtualization Engine managed node installation images
Custom

Figure 6-19 Install package selection

2. After selecting the components, the installer checks the system requirements and asks you to approve the license agreement.
3. You are asked for the installation directory. We use the same directory as for the Virtualization Engine console.
4. The installation summary is displayed, and the installation is complete.

6.4.3 Configuring the EWLM Domain Manager

You need to perform some basic configuration after the installation before starting the domain manager.

EWLM is installed in the *VE install Directory*/EWLM directory. EWLM provides some useful utilities to configure, display and change the configuration. These commands are installed in the *VE install Directory*/EWLM/bin directory:

- The **`configWizard.sh`** command provides graphical wizard to configure the domain manager.
- The **`createDM.sh`** command is also provided as a command line interface to configure the domain manager. Example 6-7 shows the command syntax of the **`createDM.sh`** command.

Example 6-7 createDM command syntax

```
Usage:
   createDM workingDir -adminUser userid -adminPW password
                       -wasPorts port -jp port
                       -ma address -mp port -dn domainName
                       -auth [None : ServerSSL : ClientServerSSL]
                       [-sslks path -sslpw password]
```

We used **`configWizardDM.sh`** to create the domain manager.

To start **configWizard**, change to the *VE install directory*/EWLM/bin directory and execute the command. Then, do the following:

1. Create a new domain manager and configure the EWLM Control Center. Select **Configure domain manager and EWLM Control Center** as shown in Figure 6-20.

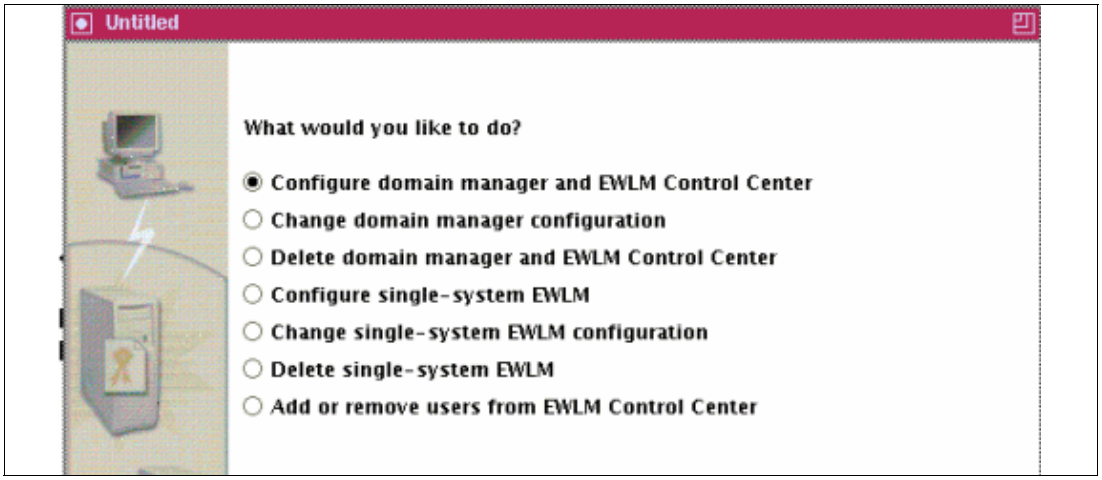
Untitled
What would you like to do?
Configure domain manager and EWLM Control Center
Change domain manager configuration
Delete domain manager and EWLM Control Center
Configure single-system EWLM
Change single-system EWLM configuration
Delete single-system EWLM
Add or remove users from EWLM Control Center

Figure 6-20 The configWizard panel

2. Specify the working directory. The working directory is created through the configWizard (with the **createDM** command).
3. Specify the domain name, the Domain Manager host address, and WebSphere Application Server ports (which hosts the Control Center).
4. Define the security level of communications between the domain manager and the managed servers. In our Environment, we select **None**.
5. Specify the SSL certificate used to communicate between the browser and the Control Center. In our ITSO environment, we use EWLM defined certificate. If you have your own certificate, select Administrator Defined and enter the keystore path and the password associated with the key.
6. If you need to create a firewall broker, specify the port of the firewall here. In our environment, we do not use firewall broker.
7. Specify the WebSphere Administrator user name that you created previously. In our environment, user name is `ewlmadm` and password is `password`.
8. Click **Finish** and the configuration begins.

In order to display the current configuration, the **displayDM.sh** command is provided. Example 6-8 shows the configuration for the domain manager in our environment.

Example 6-8 displayDM.sh command output

```
# ./displayDM.sh /ewlmwork/work

Processing displayDM request.  Please be patient as this may take a while...

WLMConfig - configurable property settings:
   ViaProxyPort/vp(null)
   TracePlugin/tlog(Off)
   InterBrokerPort/dp(null)
   InterBrokerAddress/da(null)
   JmxPort/jp(17001)
   FirewallBrokerList/fb(null)
```

```
   ReportingTrace/rt(250)
   ViaProxyHost/va(null)
   DomainName/dn(itso.ibm.com)
   JniTrace/jt(250)
   SSLKeystore/sslks(null)
   ComponentTrace/ct(250)
   DomainManagerPort/mp(17000)
   MessageLog/ml(250)
   CommunicationTrace/nt(250)
   TraceDistHubBroker/tcomm(0)
   SocksPort/sp(null)
   FirewallBrokerPort/fp(null)
   FailureLimit/fl(50)
   SSLKeystorePassword/sslpw(password suppressed)
   LoadBalancingTrace/lbt(250)
   DomainManagerAddress/ma(veaix41.itso.ibm.com)
   AuthorityLevel/auth(None)
   ProcessMode/mode(DomainManager)
   LBPublicPort/lbp(Off)
   LBSecurePort/lbs(Off)
   EventTrace/et(250)
   TraceLevel/tl(Min)
   TestComponent/t(null)
   DumpRetentionQuantity/dpn(25)
   DumpRetentionAge/dpa(30)
   SocksHost/sa(null)
WLMConfig - non-configurable property settings:
   ManagedServerFailureTime(null)
   ManagedServerIdentity(e679dec8c68853f6f712b3fd393ddf71)
   ManagedServerId(-1)
   StatisticsInterval(10)
   DomainManagerIdentity(null)
PROCESSING COMPLETE
```

After you configure the domain manager, you need to add a user to the Control Center. To add a user to the Control Center, use the wizard or the **changeCC.sh** command. We used the **changeCC.sh** command to modify the configuration of the Control Center. The EWLM control center has three different user role levels:

- The administrator can access all the functions.
- The operator has access to the managing and the monitoring tasks.
- The monitor has only access to the monitoring tasks.

Example 6-9 shows how to add the user `ewlmadm` as an administrator of the Control Center.

Example 6-9 changeCC.sh command execution example

```
# cd /opt/IBM/VE/EWLM/bin
# ./changeCC.sh -addUser /ewlmwork/work -adminUser root -adminPW password -roleUser ewlmadm
-role Administrator
```

You can check the user definition of the Control Center by executing the **displayCC.sh** command. Example 6-10 on page 143 shows the output of this command. If you want to

display the ports used by the Control Center, you can use the **-ports** option instead of the **-users** option. If you want to display groups, you can use the **-groups** parameter.

Example 6-10 displayCC.sh command output - users

```
# ./displayCC.sh -users /ewlmwork/work -adminUser ewlmadm -adminPW password
Processing displayCC -users request.  Please be patient as this may
take a while...

...processing 33% complete
...processing 66% complete
Role:  Administrator
Mapped Users:  ewlmadm
Role:  Operator
Mapped Users:
Role:  Monitor
Mapped Users:
PROCESSING COMPLETE
```

The installation and the configuration of the EWLM Domain Manager and the Control Center is now finished.

6.4.4 Installing the EWLM Managed Server

The managed server component of EWLM requires the bos.net.ewlm.rte fileset on AIX. This fileset contains the network load balancing kernel extensions for AIX. This fileset is provided with the AIX V5.2 ML03 or AIX V5.3 base operating system.

Tip: Although the installation of the managed server will succeed without this fileset, you have to install it before configuring the managed servers.

In order to confirm whether this fileset is installed on your system, you can use the **lslpp** command as shown in Example 6-11.

Example 6-11 An lslpp command output example

```
# lslpp -L |grep ewlm
  bos.net.ewlm.rte          5.3.0.10    C     F    netWLM
#
```

After you have installed the EWLM fileset, you should apply the latest fixes. You can find the latest fixes at:

http://www-912.ibm.com/eserver/support/fixes/fcgui.jsp

On the top page, select **IBM @server p5, pSeries, OpenPower, and BladeCenter HS20** as server and **AIX operating system, cluster software and Java** as Product or fix type. On the next page, click the appropriate link for your OS version in Specific fixes. Then, search by *Fileset or PTF number* with search string *bos.net.ewlm.rte*.

Tip: If you are planning to install EWLM Managed Server on AIX V5.3, we recommend the kernel mode be 64-bit configured.

If you select to install Virtualization Engine managed server installation images when you install the EWLM Domain Manager, the installation package for managed servers is located in install_directory/ManagedNode on the domain manager node. The name of the package for AIX is EWLMAixMS.bin.

- You will need to transfer this image to your managed server. If you use FTP to transfer the file, insure that the permission bits are set correctly, using the **chmod** command.
- To start the installation wizard, from the directory that contains the downloaded file, execute the following command:

  ```
  # ./EWLMAixMS.bin
  ```

  This brings up the installation wizard. Then follow the instruction of the installation wizard. The installer asks you the installation directory name and which features you wish to install. Our default installation directory is /opt/IBM/VE/EWLMMS. This value can be changed. We choose to install only the managed servers because we do not have a firewall in our environment. After the installation, you can confirm the installed package with the **lslpp** command. Example 6-12 shows an **lslpp** command display after the installation has been successful.

Example 6-12 The lslpp output after EWLM Managed Server installation

```
# lslpp -L |grep ewlm
  bos.net.ewlm.rte          5.3.0.10    C     F netWLM
  ewlm.common                1.1.0.0    C     F EWLM Common AIX Support
#
```

6.4.5 Configuring the EWLM Managed Servers

Following the installation, you need to configure the managed servers.

To configure a managed server, the **configWizardMS.sh** and **createMS.sh** commands are provided. The configuration wizard requires that you have an Xwindow installed on your system. If you do not have X Window System installed, you can configure the managed servers using the **createMS.sh** command. In our project, we used the **configWizardMS.sh** command. The wizard asks us the host name and port of domain manager, SSL setting, and firewall broker setting.

6.4.6 Verifying the installation

To verify the installation, do the following:

1. Start the Domain Manager on your management server. If you installed EWLM in the /opt/IBM/VE directory (default), the command to start the domain manager is:

   ```
   # cd /opt/IBM/VE/EWLM/bin/
   # ./startDM.sh <working directory>
   ```

   If you start the domain manager successfully, a message similar to that shown in Example 6-13 on page 145 is displayed.

Tip: Do not close your terminal window which executes the command. It is not supported to run domain manager as a background process. Thus, if you close the terminal from which you started the domain manager, the domain manager process is stopped. On AIX, you can run a process as a background process by executing command with an ampersand (&), but the EWLM Domain Manager does not support this.

Example 6-13 startDM.sh example

```
# ./startDM.sh /ewlmwork/work
Starting EWLM Domain Manager...

A message is displayed when the EWLM Domain Manager completes activation.
WARNING: Closing this message window ends the EWLM Domain Manager.

<9:51:32.5132> EWLM successfully started
```

2. Start the Control Center on your management server. If you installed EWLM in the /opt/IBM/VE directory (default), the command to start the domain manager is:

   ```
   # cd /opt/IBM/VE/EWLM/bin/
   # ./startWAS.sh <working directory>
   ```

   This command is not an interactive command. This command starts a WebSphere Application Server. The WebSphere Application Server runs as a daemon.

3. Start the Managed Server on your managed server. In order to start a managed server on AIX, you need to enable the EWLM service. Execute **`smit ewlm`** and select **Change/Show status of EWLM services** and **Enable the EWLM system**.

   ```
   # cd /opt/IBM/VE/EWLMMS/bin
   # ./startMS.sh <working directory>
   ```

   If you start the managed servers successfully, the message as shown in Example 6-14 is displayed. Do not close your terminal window which executes this command. It is not supported to run EWLM Managed Server as a background process, either.

Example 6-14 startMS.sh example

```
# ./startMS.sh /ewlmwork/work
Starting EWLM Managed Server....
A message is displayed when the EWLM Managed Server completes activation.

WARNING: Closing this message window ends the EWLM Managed Server.

<9:51:36.5136> EWLM successfully started
```

4. Log in to the Control Center and to check the status of the domain manager and the managed servers. To connect to the Control Center, you need to know the host name and the port number of the Control Center. The **`displayCC.sh`** command is provided to show ports for Control Center. Example 6-15 shows the output of **`displayCC.sh`**.

Example 6-15 displayCC.sh command output - ports

```
# ./displayCC.sh -ports /ewlmwork/work -adminUser ewlmadm -adminPW password
Processing displayCC -ports request.  Please be patient as this may
take a while...

...processing 33% complete
...processing 66% complete

 ...Ports assigned to EWLM Control Center:
      - HTTP    17103
      - HTTPS   17104
...Use -changeCC -controlCenterPorts to change these ports if desired.

...Ports assigned to WebSphere Admin Console
      - HTTP    17101
      - HTTPS   17100
...Use -changeCC -adminConsolePorts to change these ports if desired.
```

```
...Port assigned to WebSphere Admin: 17109
...Use -changeCC -adminPort to change this port if desired.

PROCESSING COMPLETE
```

We can find that port 17103 is assigned for HTTP connection. In our environment, the URL is:

```
http://veaix41.itso.ibm.com:17103/webui/
```

If you access to the HTTP port, you are redirected to the HTTPS port. Figure 6-21 shows the welcome page of EWLM Control Center. Enter the user name and password that you added by the **changCC.sh** command.

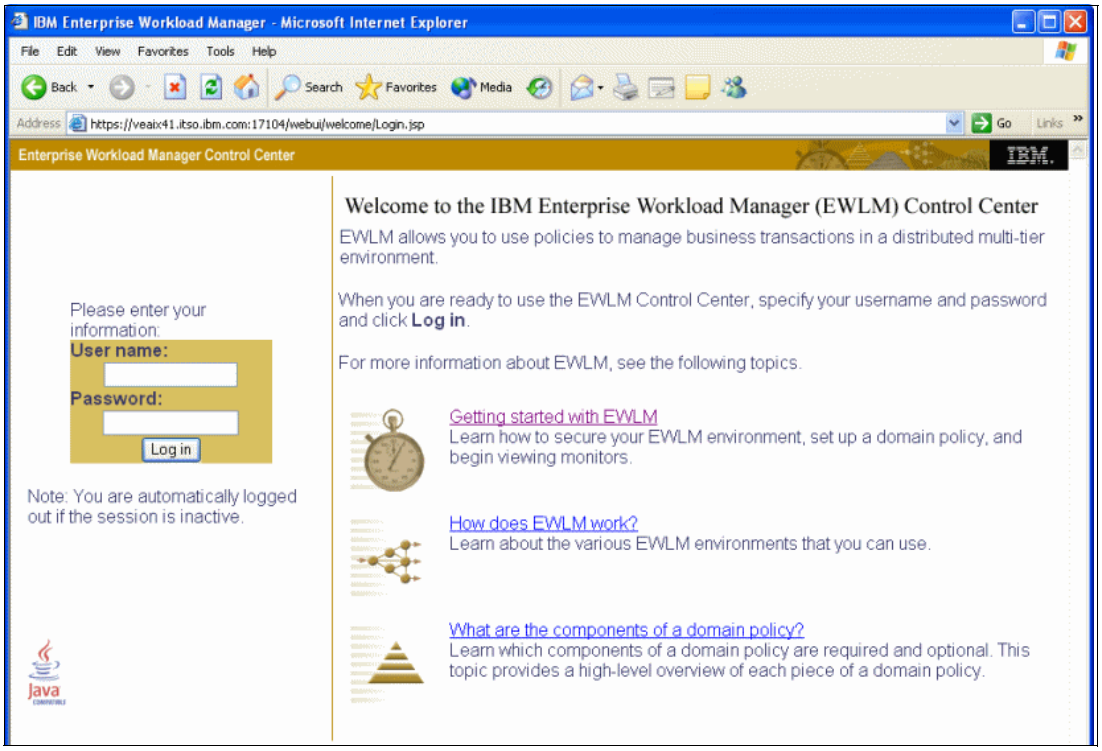
IBM Enterprise Workload Manager - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address https://veaix41.itso.ibm.com:17104/webui/welcome/Login.jsp
Go Links
Enterprise Workload Manager Control Center
IBM
Welcome to the IBM Enterprise Workload Manager (EWLM) Control Center
EWLM allows you to use policies to manage business transactions in a distributed multi-tier environment.
When you are ready to use the EWLM Control Center, specify your username and password and click Log in.
For more information about EWLM, see the following topics.
Please enter your information:
User name:
Password:
Log in
Note: You are automatically logged out if the session is inactive.
Getting started with EWLM
Learn how to secure your EWLM environment, set up a domain policy, and begin viewing monitors.
How does EWLM work?
Learn about the various EWLM environments that you can use.
What are the components of a domain policy?
Learn which components of a domain policy are required and optional. This topic provides a high-level overview of each piece of a domain policy.
Java

Figure 6-21 Welcome page of the Control Center

If you can log in to the EWLM Control Center, your user has been configured successfully.

5. Click **Managed Servers** on the left navigation panel. If the connection between the domain manager and the managed servers is successfully established, the managed server host name is displayed on the right panel and its status is shown as Active, as shown in Figure 6-22 on page 147.

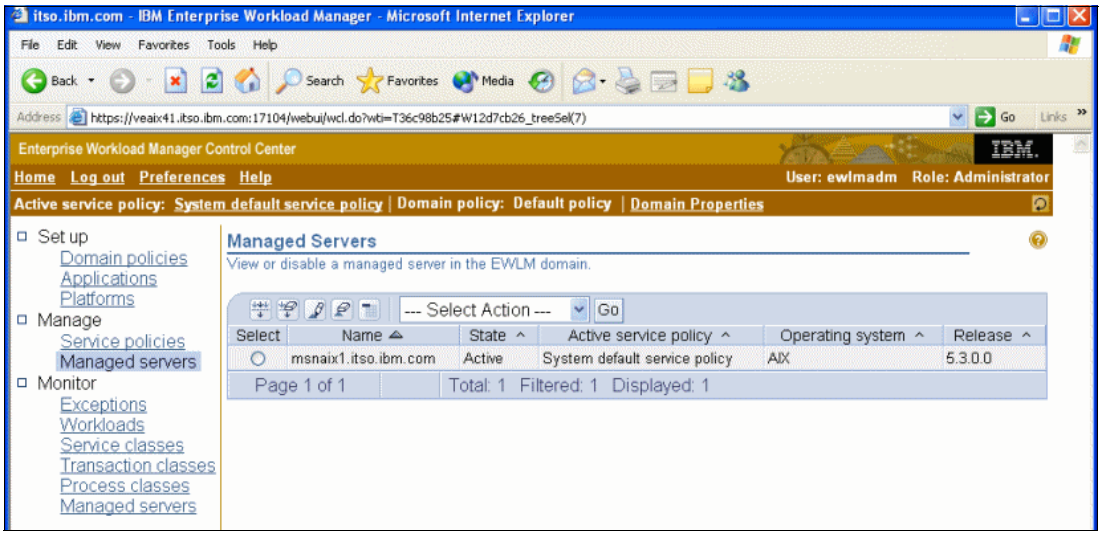

Figure 6-22 Managed Server view

6. You can verify that instrumentation has been enabled using the **`lsarm`** command:

```
$ lsarm -a
```

The command might print the following ARM registered applications:

```
APPL: IBM DB2 Universal Database
APPL: WebSphere
APPL: IBM Webserving Plugin
APPL: Apache HTTP Server
```

In case that domain manager failed to communicate with managed server, no server is displayed or the server status is `communication error`. If the Control Center cannot connect to domain manager, you get an error message in the window.

6.5 Installing Tivoli Provisioning Manager on top of the Virtualization Engine components on a pSeries with AIX

This section describes how to install Tivoli Provisioning Manager on an AIX management server, on top of the Virtualization Engine components (Virtualization Engine console, IBM Director, EWLM). We already have management servers of Virtualization Engine console, EWLM, and IBM Director and managed servers of IBM Director and EWLM. In this section, we add the provisioning component, Tivoli Provisioning Manager, in our environment.

Figure 6-23 shows the logical topology used in this chapter.

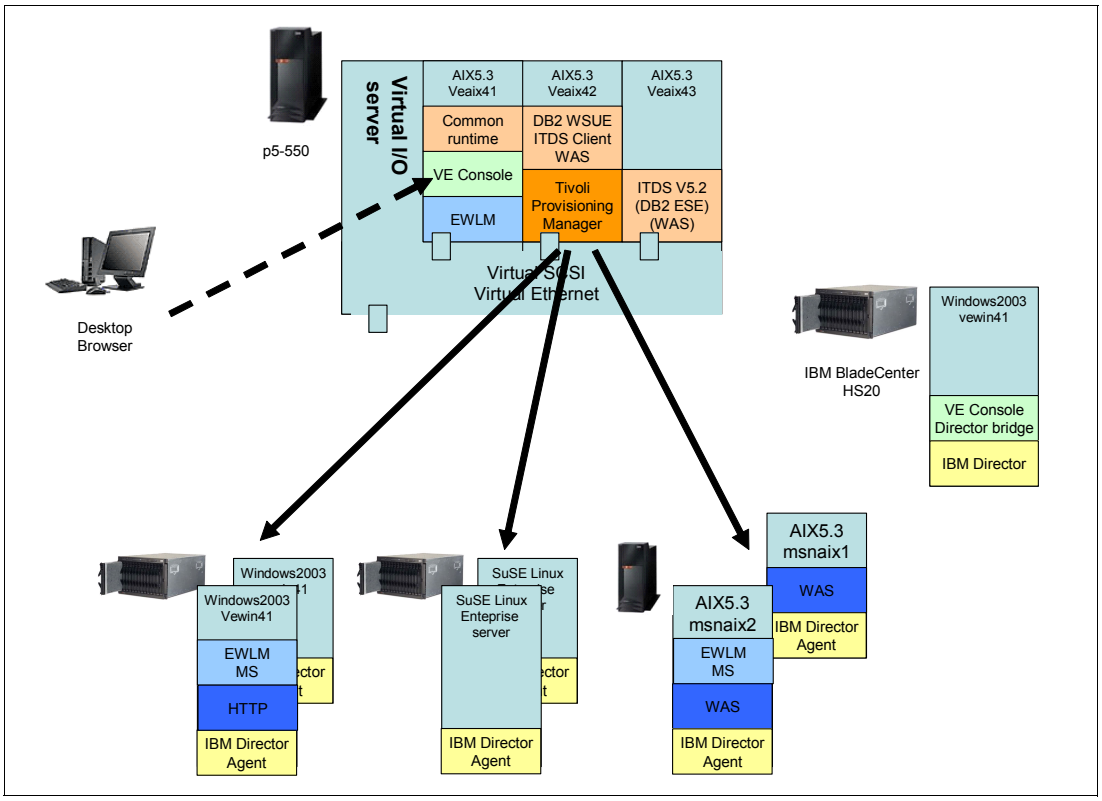

Figure 6-23 Topology

We use two LPARs running AIX V5.3 on a pServer p5-550 as management servers for Tivoli Provisioning Manager.

On AIX, Tivoli Directory Server for Tivoli Provisioning Manager must be installed separately. In our environment, we already have Tivoli Directory Server V5.1 for Virtualization Engine console. However, Tivoli Provisioning Manager does not support to connect to Tivoli Directory Server V5.1. Therefore, we have another LPAR running Tivoli Directory Server V5.2. We install Tivoli Directory Server V5.2 on AIX, but it is possible to install Tivoli Directory Server on another platform. For more information about the supported Tivoli Provisioning Manager Topology, refer to *Tivoli Provisioning Manager Installation Guide for UNIX*, GC32-1615.

Our Tivoli Provisioning Manager and Tivoli Directory Server server are running AIX V5.3 64-bit kernel. At the time of writing, Tivoli Provisioning Manager only supports AIX V5.2 64-bit or V5.3 64-bit running on POWER5 system. You can check the latest information at:

http://www-306.ibm.com/software/tivoli/products/prov-mgr/requirements.html

Tivoli Provisioning Manager requires several prerequisites software. Table 6-3 shows the software stack of our servers.

Table 6-3 Software stack

Host name	Software
veaix42	Tivoli Provisioning Manager V2.1 with FP02
	DB2 Universal Database Workgroup Server Unlimited Edition V8.1
	WebSphere Application Server V5.1 with cumulative Fix 3
	Tivoli Directory Server V5.2 client
	OpenSSH server/client (See 6.5.8, “Steps 7 to 9: Post-installation tasks” on page 158)
	Prerequisites utilities (See 6.5.3, “Step 2: Installing prerequisites utilities” on page 150)
veaix43	Tivoli Directory Server V5.2 Server/Client
	DB2 Universal Database Enterprise Server Edition V8.1

6.5.1 The installation process

The installation process of Tivoli Provisioning Manager on AIX is described in *Tivoli Provisioning Manager Installation Guide for UNIX*, GC32-1615. Refer to this document for detailed information.

In this chapter, we describe the installation process as we did it in the our environment, with additional information when required.

Overview of the installation process

Tivoli Provisioning Manager requires several prerequisites packages and the configuration of specific users and file systems. After we configure them, we install WebSphere Application Server, DB2 Universal Database, and Tivoli Directory Server server and Client. Our environment is a two node configuration. We install WebSphere Application Server, DB2 Universal Database and Tivoli Directory Server client on a same machine and the TDS server on another machine.

The installation of the Tivoli Provisioning Manager product requires the following steps:

1. Installation and configuration of systems prerequisites, such as WebSphere and WebSphere MQ.
2. Installation of prerequisite utilities, which are open source packages.
3. Installation of Tivoli Directory Server (which needs a WebSphere Application Server installed) and DB2 as the database.
4. Installation of prerequisites software on the server node.
5. Installation of patches for Tivoli Directory Server and WebSphere.
6. Installation of the Tivoli Provisioning Manager code.
7. Apply the latest fix pack on Tivoli Provisioning Manager.

8. Import a DCM file.
9. Configuration of SSH.
10. Installation verification.

These steps are detailed in the following sections.

6.5.2 Step 1: Configuring system prerequisites

Tivoli Provisioning Manager requires WebSphere Application Server and WebSphere embedded messaging. Before we start the installation, we need to create a user group and a file system for WebSphere embedded messaging.

This configuration must be done on the node where the WebSphere Application Server is installed. For example, veaix42 is the node for WebSphere Application Server in our environment. We create /var/mqm file system and user mqm and group mqm and mqbrkr for WebSphere embedded messaging. These user name and file system name are fixed. The detailed steps are described in Chapter 5 in *Tivoli Provisioning Manager Installation Guide for UNIX*, GC32-1615.

6.5.3 Step 2: Installing prerequisites utilities

Tivoli Provisioning Manager requires several open source packages. Table 6-4 shows the list of these packages for which you need to use the openSSH package.

OpenSSH is included in the AIX Base Operating System CD-ROM for AIX V5.3. If you are planning to use AIX V5.2, you need to download it from:

`https://www14.software.ibm.com/webapp/iwm/web/preLogin.do?source=aixtbx`

You can download the other opensource software from the AIX Toolbox for Linux Application Web site:

`http://www-1.ibm.com/servers/aix/products/aixos/linux/download.html`

Table 6-4 Utilities

Fileset or package name	Obtained from
xlC.rte 6.0 or higher	AIX Base Operating System CD-ROM or download from fix central (for AIX5L 5.2)
Expect 5.32 or higher	AIX Toolbox for Linux Applications Web page
bash	AIX Toolbox for Linux Applications Web page
bash-doc	AIX Toolbox for Linux Applications Web page
unzip	AIX Toolbox for Linux Applications Web page
tcl	AIX Toolbox for Linux Applications Web page
tk	AIX Toolbox for Linux Applications Web page

These are provided as an rpm package. Therefore, you need to use the **rpm** command to install them. For example, if you install the Expect 5.32 package, execute the following command.

```
# rpm -ivh expect-5.34-8.aix4.3.ppc.rpm
```

When you install the opensource software, you need to look at the dependencies. We installed the packages in the following order:

1. tcl
2. tk
3. expect
4. bash
5. unzip

6.5.4 Step 3: Installing and configuring the Tivoli Directory Server V5.2 code

In an AIX environment, we cannot install the Tivoli Directory Server V5.2 and Tivoli Provisioning Manager on the same box because Tivoli Provisioning Manager installs DB2 Universal Database Workgroup Edition V8.1. Therefore, we need to install Tivoli Directory Server separately. In order to install Tivoli Directory Server, we need to install WebSphere Application Server; WebSphere Application Server5.1 for AIX Tivoli Directory Server V5.2 is provided with Tivoli Provisioning Manager V2.1.

Note: In this project, we used WebSphere Application Server Network Deployment 5.2 Tivoli Directory Server V5.2.

It is possible to install IBM Directory Server by using SMIT or the **installp** command. However, in our installation we used the graphical installer. The installation process is described in the*IBM Tivoli Directory Server Installation and Configuration Guide*, SC32-1338. Refer to this document for the detailed installation using SMIT or **installp**.

Before you start the installation, you have to verify the following:

- The system requirements.
- The /tmp file system must have enough free space to expand a DB2 image. If you do not have enough space, the installation will fail. In our environment, it is 300 MB.
- The status of an asynchronous I/O device. An asynchronous I/O device must be available to finish the installation. If its status is defined, execute **smitty aio** and select **Configure Defined Asynchronous I/O**. In order to make aio available automatically when the system is restarted, you have to execute the **smitty chgaio** command and change the STATE to be configured at system restart parameter to available as shown in Figure 6-24 on page 152.

```
Change / Show Characteristics of Asynchronous I/O

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                                [Entry Fields]
  MINIMUM number of servers                    [1]                          #
  MAXIMUM number of servers per cpu            [10]                         #
  Maximum number of REQUESTS                   [4096]                       #
  Server PRIORITY                              [39]                         #
  STATE to be configured at system restart      available                   +
  State of fast path                            enable                      +
```

Figure 6-24 SMIT example of asynchronous I/O setting

To start the installation, put the CD-ROM in your machine and mount it. Then execute the **setup** command. This command launches the graphical installer. After the installer is started, do the following:

1. Click **Next** and accept the license agreement and click **Next**.
2. Select your language and click **Next**. In our environment, we select English.
3. The summary of the configuration is displayed. Confirm the settings and click **Next** to start the installation.
4. Select all the packages. If you have installed any package in the list, the Installation guide provides special consideration. Refer to the Installation Document.
5. The verification window is displayed with your configuration. Click **Next** to continue. The installer start to copy the files.
6. After the file copy is finished, several readme files are displayed. Click **Next**.
7. Click **Finish**.

When you click finish, the **ldapxcfg** command is started automatically. We need to configure an administrator DN and a database using this command.

1. In order to create an administrator DN, click **Administrator DN/password**, and enter your administrator DN and password. Then, click **Next**. In our environment, the administrator DN is cn=root.
2. Next, we create a database for the LDAP serer. Before we create the database, we need to create an owner user of the database. We create the user db2isnt1 as the owner of the database. This user must satisfy the following criteria.
   - The user must be a member of the group that created during DB2 installation. Normally, this group name is dbsysadm on AIX. The user's home directory is owned by this group.
   - The password must be set correctly and ready to use. By default, a user is asked to input a new password on the fist time to log in to the system. In order to avoid this situation, you need to execute the following command.

     ```
     # pwdadm -f NOCHECK db2inst1
     ```

   In order to create a database, click **Configure database**.

   a. If a user ID and password are requested, type the user name and password you have created.

   b. If the database name is required, type a name from 1 to 8 characters.

c. If the database location is requested, type the location for the database in the Database location field. This must be a directory.

Tip: Ensure that the DB2 database directory have enough space to create new database. If that has not enough free space, the create database operation fails.

d. If a character set selection is requested, click the type of database that you want to create. If you want to use language tags, the database must be a UTTF-8 database.

e. The verification window is displayed with your configuration. Click back if you want to change them. To begin the configuration process, click **Next**.

Now, the LDAP server is configured. However, the LDAP cannot start on boot time by default. If you want to start the LDAP server at boot time, you need to add an entry in the /etc/inittab file with the following command:

```
# /usr/sbin/mkitab "ldap:2:once:/usr/ldap/bin/slapd -f /etc/ibmldap.cfg > /dev/null 2&>1"
```

In order to check the installation, we recommend to reboot the system.

6.5.5 Step 4: Installing the prerequisite software on the server node

Tivoli Provisioning Manager provides a prerequisite installer for the prerequisite software. This installer provides a method for installing IBM middleware and its related components. WebSphere Application Server, DB2 Universal Database Workgroup Server and Tivoli Directory Server client are installed by this installer.

Before installation, verify the following:

- Make sure you install the component node you want. In our two node topology, WebSphere Application Server, DB2 Universal database and Tivoli Directory Server client are installed on a same machine.
- The prerequisites configuration has been done on the WebSphere Application Server node.
- The prerequisites utilities are installed on the Tivoli Provisioning Manager server node.

In order to execute the installation, the following steps are required.

1. If you have a CD-ROM, insert your CD-ROM and mount it. If you use a tar image, expand it in your local disk.
2. Start the installer using the following command:

   ```
   # <mount point>/setupaix.bin
   ```

   Attention: Do not change the directory to mount point. Changing directory to mount point will lock the CD-ROM drive. You need to swap CDs during the installation process.

3. If a language selection is required, select a language and click **OK**. The support for this language will also be installed for the prerequisite software.
4. Specify the directory on which the copy is done. The Tivoli Software Installer copies itself to the local hard drive to free up the CD drive. While the Tivoli Software Installer copies the code to the hard drive, the installer screen might not be visible. After a few seconds, a new installer screen appears. Select language again.
5. Select the products you are going to install. In our two node topology, we select DB2 Universal Database, IBM Directory Server 5.2 with DB2 8.1, and WebSphere Application

Server V5.1. Even if you select IBM Directory Server 5.2 with DB2 8.1, only the client portion of the IBM Directory Server is installed with DB2 on AIX.

6. If we select DB2 Universal Database, the installer asks us whether we want to install the client and the server or the client only. In our two-node topology, we do not use a remote database server. Therefore, we select DB2 Workgroup Server with Fix Pack 3 here. Then, the installer asks us about a database user name and password. The default user name is db2inst1. Even if you select IBM Directory Server 5.2 with DB2, only the client portion of Tivoli Directory Server is installed. The installer provides a message.
7. Before the product installation, the installer asks us about the location of the installation code. If you are installing by using a CD, insert and mount it. Then, type the appropriate location. In our environment, the following products have been installed:

   a. DB2 Universal Database, Workgroup Server Unlimited Edition, Version 8.1
   b. DB2 Universal Database, Workgroup Server Unlimited Edition, Version 8.1 Fix Pack 3
   c. IBM Tivoli Directory Server V5.2
   d. WebSphere Application Server 5.1

   During the product installation, the progress bar appears. This progress bar is updated only when each product installation is finished.

Attention: If you have installed WebSphere Application Server on an AIX 64–bit machine, which has its code page set to a language other than en_US, you must set the environment variables. Without these environment variables set, you might experience a problem with not all of the expected WebSphere file sets being listed when issuing the **`lslpp`** command. This is a documented WebSphere Application Server issue that might be already corrected when you read this redbook.

- Set the EXTSHM environment variable by entering the following commands:

```
EXTSHM=ON
export EXTSHM
```

- Set the LDR_CNTRL environment variable by entering the following commands:

```
LDR_CNTRL
export LDR_CNTRL=MAXDATA=0X30000000
```

Now, the product installation is finished. You need to apply the mandatory patches to the Tivoli Directory Server and the WebSphere Application Server.

6.5.6 Step 5: Applying mandatory patches

After the installation of the prerequisite software, you need to apply the following patches to the Tivoli Directory Server, WebSphere Application Server, and WebSphere MQ V5.3 fix. These fix packs are located on the Tivoli Provisioning Manager and Tivoli Intelligent Orchestrator Version 2.1 Generic Fixes CDs.

To apply Tivoli Directory Server fixes, do the following steps:

1. Logon to the LDAP server node as root. In our environment, that is veaix43.itso.ibm.com.
2. If Tivoli Directory Server is running, stop it by executing the following command.

```
# /usr/ldap/bin/ibmdirctl -D"cn=root" -w? stop
```

3. Change `cn=root` to your LDAP administrator domain name.
4. From the /ITDS directory on the CD, copy the v3.ibm.at file to the ITDS_installdir/etc directory. ITDS_installdir is the location of the Tivoli Directory Server installation directory.

5. Restart the Tivoli Directory Server by executing the following command:

```
# /usr/ldap/bin/ibmdirctl -D"cn=root" -w? start
```

To apply WebSphere Application Server 5.1 cumulative fix 3, do the following steps:

1. Log on to the WebSphere Application Server host as root.
2. Stop the WebSphere Application Server by executing the following command.

```
# WAS_installdir/bin/stopServer.sh server1
```

3. Switch to the WebSphere Application Server installation directory and create a subdirectory called WAS_installdir/update/was510_cf3.
4. Copy the Cumulative Fix 3 file from the WAS/CF3 directory on the CD to the WAS_isntalldir/update/was510_cf3 directory and unzip it. For AIX, the file name is was510_cf3_aix.zip. To expand the file, execute the following command:

```
# unzip was510_cf3_aix.zip
```

5. Run the UpdateWizard with the command **updateWizard.sh**. During the installation, ensure that the Install fix packs option is selected.
6. After installing fix pack, restart WebSphere Application Server.

To apply the WebSphere MQ fix:

1. Logon to the WebSphere Server host as root.
2. Copy the file IC38409.unix.tar.Z from the WAS/IC38409 directory on the CD to your local drive. Expand the file with the following command:

```
# zcat IC38409.tar.Z | tar -xvf -
```

3. Stop the all running queue managers and channel listeners.
4. Confirm the permission and ownership of com.ibm.mqjms.jar and com.ibm.mq.jar files located on /usr/mqm/java/lib directory and write them down.
5. Copy the new versions on top of the existing com.ibm.mqjms.jar and com.ibm.mq.jar files. Make sure that the permission and ownership is same as old ones.
6. Restart the queue managers and channel listeners.

6.5.7 Steps 6: Installing the Tivoli Provisioning Manager code

After having installed the prerequisites products and applied the mandatory patches, we can start the installation of Tivoli Provisioning Manager. Tivoli Provisioning Manager provides the following installation methods:

- The graphical install shield.
- The silent install, using a response file.

  We can install Tivoli Provisioning Manager by using a response file that is created during the installation or template provided with Tivoli Provisioning Manager. To create a response file, execute the **setupaix.bin** command with **-options-record response_file_name**.

  For more information, refer to Chapter 4 in *Tivoli Provisioning Manager Installation Guide for UNIX*, GC32-1615. This method of installing is recommended for advanced users.

- The no-graphical installation, or the console mode installation.

  This method is useful if your machine does not have a graphic card. The installation process is similar to the one using the graphical installer. In order to run a non-graphical

installer, execute **setupaix.bin** with the **-console** option. For more information, refer to Chapter 4 in *Tivoli Provisioning Manager Installation Guide for UNIX*, GC32-1615.

In this chapter, we describe the installation method using the graphical install shield. Before starting the installation, verify the following.

- The hardware and software requirements.

  If you are running AIX V5.3 on your Tivoli Provisioning Manager server node, you need to apply Tivoli Provisioning Manager Fix Pack 2.1.0.1-TIV-TPM-FP0001 after the installation is completed.
- The installation of the prerequisite products is completed and the mandatory fixes are applied.
- Ensure which node you want to install Tivoli Provisioning Manager server. In our environment, `veaix42.itso.ibm.com` is the node for Tivoli Provisioning Manager server.
- Ensure the database server and the LDAP server are running correctly.
- Ensure the AIX Web based System Management is disabled. Otherwise, WebSphere Application Server startup will fail. If it is enabled, execute the following command:

  ```
  #/usr/websm/wsmserver -disable
  ```
- Ensure that the JMS server initial state is set to started. In order to confirm it, do the following:
  a. Start WebSphere Application Server Admin Console.
  b. Go to Server → Application server → Server1 → Server components → JMS server → Initial state → Started.
  c. Save the change.
  d. Stop WebSphere Application Server.
- Create the tioadmin user which uses bash as its login shell.

Tip: If you are planning to install on AIX V5.3, the update installer is required in order to install the Base Tivoli Provisioning Manager v 2.1.0 on AIX 5.3. Download the new installer called “setup.jar” from Tivoli Provisioning Manager Fix Pack 2.1.0.1-TIV-TPM-FP0001 Web site.

You need to replace “setup.jar” by the new one. Therefore, if you are planning to install on AIX V5.3, copy the Tivoli Provisioning Manager Version 2.1 CD to your local drive. Then, replace <copy_dir>/setup.jar by the new one you downloaded from the Web site.

In order to install the Tivoli Provisioning Manager code, do the following:

1. Insert the Tivoli Provisioning Manger Version 2.1 CD into the CD drive and mount the CD-ROM drive, but do not change the directory by the mount point. Launch the installer by entering the following:

   ```
   mount_point/setupaix.bin
   ```
2. If the language selection is required, select the language you want to use. Click OK and accept the license agreement.
3. The installer verifies if the system meets the prerequisites. This might take a few minutes.
4. If the Tivoli Provisioning Manager Installation Directory is displayed, accept the default installation path or enter a custom directory path. The default installation path is /opt/IBM/tivoli/.

5. If the administrator user account and user groups are required, enter and verify the password for the tioadmin.
6. If the language selection is required, select the languages you want to install.
7. On the database configuration panel, type your configuration. Table 6-5 shows our database configuration. If you use the prerequisites installer to install DB2 Universal Database, the host port is 50000. If you manually configure the database server, select **Do not perform Database server configuration steps**.

Table 6-5 Database server configuration

Entry	Value
Fully Qualified Host Name or IP Address of the IBM DB2 Universal Database	veaix42.itso.ibm.com
Host Port	50000
Database Server Instance Owner	db2inst1
Database Server Instance Owner password	password
Database User name	db2inst1
Database User Password	password
Load DB2 instance SQLLIB directory	/home/db2inst1

8. On the LDAP server configuration panel, type your configuration. Table 6-6 shows our LDAP server configuration. If you wish to manually configure the directory server, check the “Do not perform Directory server configuration steps”.

Table 6-6 LDAP configuration

Entry	Value
Base domain name	dc=ibm, dc=com (default)
LDAP Administrator User Name	cn=root
LDAP Administrator User Password	password
Fully Qualified Host Name or IP Address	veaix43.itso.ibm.com
Host Port	389
LDAP Client Installation Directory	/usr/ldap

9. On the WebSphere Application Server configuration panel, enter your configuration. Table 6-7 shows our WebSphere Application Server configuration. We cannot use IP addresses.

Table 6-7 WebSphere Application Server configuration

Entry	Value
DNS Suffix Name	itso.ibm.com
Installation Directory	/opt/WebSphere/AppServer

10. The next screen provides you with the option to install Tivoli NetView. This is optional.

11. The installation preview summarizes the information you have entered during the installation. Review the information to ensure that it is accurate. If the information is correct, click **Next** to start the product installation.
12. After the installation, the summary panel is displayed. If you want to start Tivoli Provisioning Manager services now, select the check box.
13. Click **Finish** to exit the installation wizard.

Now, you have completed the installation of the Tivoli Provisioning Manager code.

6.5.8 Steps 7 to 9: Post-installation tasks

The section describes the post-installation tasks that you need to perform.

Step 7: Applying fix pack for Tivoli Provisioning Manager (mandatory for AIX V5.3)

For AIX V5.3, applying FP01 is mandatory. For information about the fix pack installation, refer to the readme file for the fix pack.

Step 8: Importing a data center model

The data center model provided with the Tivoli Provisioning Manager base CD needs to be modified. If you try to import it, you will get error. We changed the datacentermodel.xml file to import with no error.

To import the data center model, the **xmlimport.sh** command is provided. To import the data center model, do the following:

1. Log on as tioadmin.
2. Execute the following commands:

```
$ cd $TIO_HOME/xml
$ $TIO_HOME/tools/xmlimport.sh file:$TIO_HOME/xml/datacentermodel.xml
```

Step 9: Installing and Configuring SSH

OpenSSL version 3.4 or higher must be installed and configured. To do so, the following steps are required:

1. Install the OpenSSL package, which you can find at:

   https://www14.software.ibm.com/webapp/iwm/web/preLogin.do?source=aixtbxi

   Registration is required. After registering, you are redirected to a Web page where you can download OpenSSL.

> **Tip:** If your environment is AIX V5.3, you could find OpenSSL packages in your AIX Bos CD-ROM.

2. Install the following filesets from the AIX Base installation media:
   - openssh.base
   - openssh.license
   - openssh.msg.en_US
   - openssh.man.en_US

If the filesets were not found on the AIX Base installation media, you can download the from the following:

http://www-1.ibm.com/servers/aix/products/aixos/linux/download.html

3. Start the sshd daemon by running the following command:

```
# /usr/bin/startsrc -s sshd
```

Tip: If the AIX machine on which OpenSSH is installed has GSA installed, the SSH daemon will not start. You need to first check to see if the sshd user exists on the system. If not, it should be created with the following commands:

```
# mkgroup sshd
# mkuser -a pgrp=sshd login=false home=/var/empty gecos="OpenSSH privilege
separation" account_locked=true sshd
```

Some text might appear on separate lines for presentation purposes only.

4. As user tioadmin, configure SSH so that the server can communicate with relevant users on other systems and components of the data center.

Tip: Ensure that you are logged on to user ID tioadmin directly. Do not use **su - to tioadmin** or the following steps will fail.

In order to configure SSH, Run the following commands:

```
$ ssh-keygen -t rsa -N "" -f $HOME/.ssh/id_rsa
$ cat $HOME/.ssh/id_rsa.pub >> $HOME/.ssh/authorized_keys
```

You can test this by running:

```
ssh -v tioadmin@localhost
```

where *localhost* is your host name. If SSH is properly configured, you are not prompted for a password.

Copy the public key for user tioadmin to the servers that Tivoli Provisioning Manager will be managing in your data center.

It is required to configure SSH to accept connections from new hosts without prompting for confirmation. Create a file in /home/thinkcontrol/.ssh called config. The file should contain the following line:

```
StrictHostKeyChecking no
```

Copy the id_rsa.pub file into the authorized keys file of the target server administrative account to the database server and any managed server running Cygwin. This file must be copied to any new servers brought under Tivoli Provisioning Manager's control.

Setting up tioadmin user shell environment, if applicable.

The user tioadmin uses the bash shell as the login shell. This works well for a line-mode login (for example, via telnet). Log in via a window manager such as a Common Desktop Environment might not be able to create a complete login environment, which is required for Tivoli Provisioning Manager to function properly. This can be resolved by doing the following:

1. Create the .bashrc file in the tioadmin's home directory, and insert the following single line:

```
. $HOME/.profile
```

2. Save the file.

3. Edit the .dtprofile in tioadmin's home directory and uncomment the following line: DTSOURCEPROFILE=true. This file is created automatically when the user tioadmin logs in CDE for the first time.
4. Log in as user tioadmin again to the Common Desktop Environment.

6.5.9 Step 10: Verifying the installation

In order to start the installation, do the following;

1. Ensure that DB2 is started. In order to start DB2, log in to the DB2 server as instance owner (in our environment, `db2inst1`):

```
$ db2start
```

2. Ensure that the LDAP server is started. In order to start the LDAP server, log in to the LDAP server as root:

```
# /usr/bin/ibmdirctl -D"cn=root" -w? start
```

3. Start the Tivoli Provisioning Manager server. Log in to the Tivoli Provisioning Manager server as tioadmin:

```
$ cd $TIO_HOME/tools
$ ./tio.sh start
```

After you have started the server, open your Web browser, and go to the following URL, where hostname.com is your Tivoli Provisioning Manager server name:

```
http://hostname.com:9080/tcWebUI/index.jsp
```

If you can log in, the installation is successful.

7

Installing the Virtualization Engine console, IBM Director, TPM, and EWLM on Blade with Windows

This chapter discusses the key concepts when installing and configuring the following Virtualization Engine system services components:

- IBM Virtualization Engine console
- IBM Director
- IBM Enterprise Workload Manager
- IBM Tivoli Provisioning Manager

This chapter supplements the step-by-step installation instructions located in the IBM @server Software Information Center (IBM Information Center) at:

http://publib.boulder.ibm.com/infocenter/eserver/v1r1/en_US/index.htm?info/icmain.htm

Our environment is based on output from the Virtualization Engine console planning advisor, available at:

http://publib.boulder.ibm.com/infocenter/eserver/v1r1/en_US/info/veinfo/planner.htm

We ran the planning advisor for our Windows environment and it produced the output shown in Figure 7-1 on page 162. This output was the starting point of our installation.

© Copyright IBM Corp. 2005. All rights reserved.

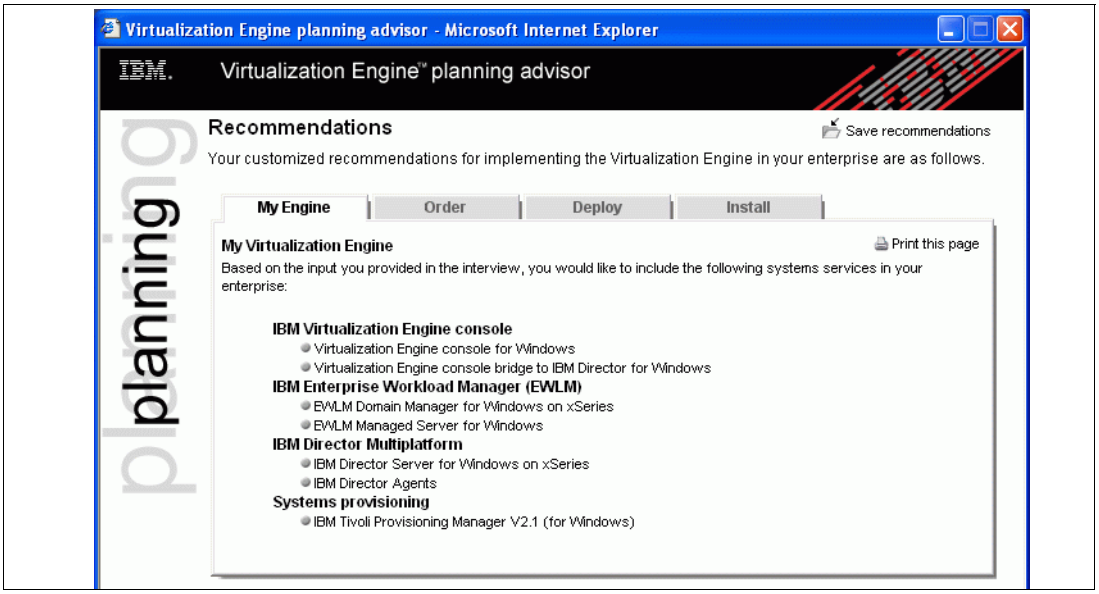
Virtualization Engine planning advisor - Microsoft Internet Explorer
IBM.
Virtualization Engine™ planning advisor
planning
Recommendations
Save recommendations
Your customized recommendations for implementing the Virtualization Engine in your enterprise are as follows.
My Engine
Order
Deploy
Install
My Virtualization Engine
Print this page
Based on the input you provided in the interview, you would like to include the following systems services in your enterprise:
IBM Virtualization Engine console
Virtualization Engine console for Windows
Virtualization Engine console bridge to IBM Director for Windows
IBM Enterprise Workload Manager (EWLM)
EWLM Domain Manager for Windows on xSeries
EWLM Managed Server for Windows
IBM Director Multiplatform
IBM Director Server for Windows on xSeries
IBM Director Agents
Systems provisioning
IBM Tivoli Provisioning Manager V2.1 (for Windows)

Figure 7-1 Virtualization Engine planning advisor output

7.1 Before the installation

Before the actual installation, we describe the topology, the installation flow, and the few pre-installation steps that we have to consider.

7.1.1 The topology

Our environment consists of three management servers and two managed servers, as illustrated in Figure 7-2.

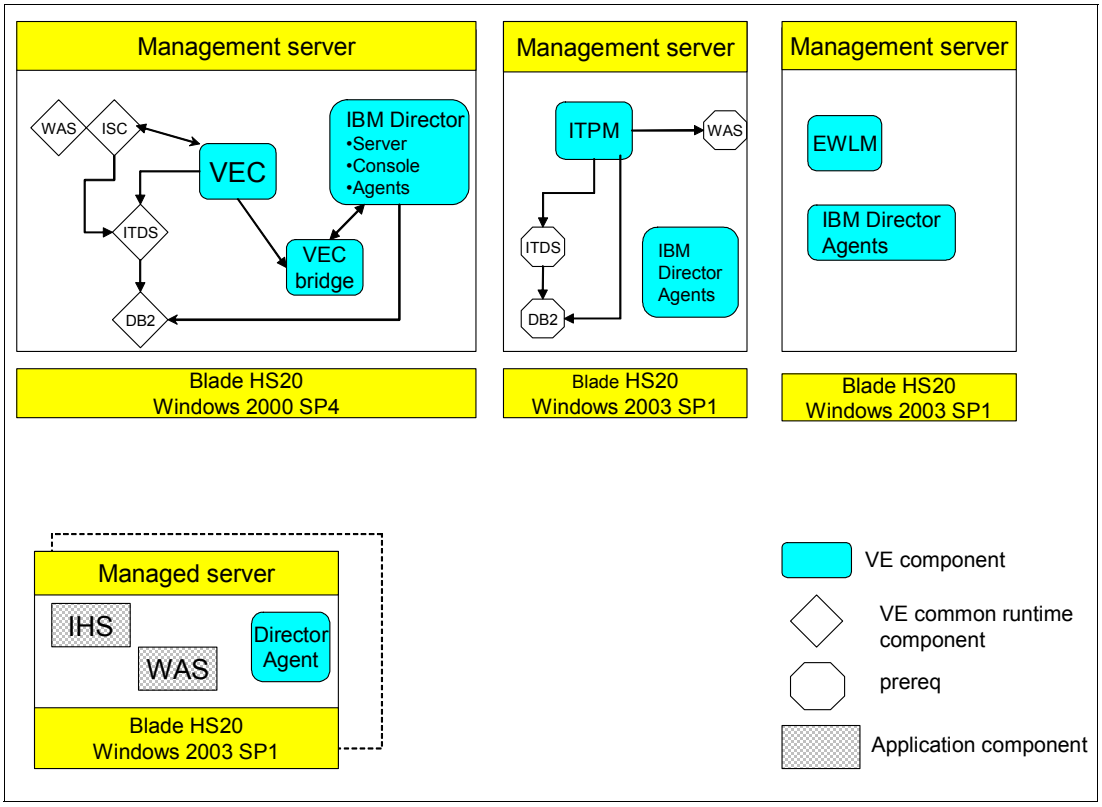

Figure 7-2 Management servers topology

The management servers have the following characteristics:

- vewin21.itso.ibm.com
  IP: 10.1.1.178
  Operating System: Windows 2000 SP4
  Hardware: BladeCenter HS20
  Software Installed: Virtualization Engine console, WebSphere Application Server 5.1, Tivoli Directory Server 5.1, DB2 8.1, Director Multiplaform Server (includes agents and console), Virtualization Engine console bridge to IBM Director
- vewin22.itso.ibm.com
  IP: 10.1.1.179
  Operating System: Windows 2003 SP1
  Hardware: BladeCenter HS20
  Software Installed: Enterprise Workload Manager, IBM Director Agent

- vewin23.itso.ibm.com
  IP: 10.1.1.180
  Operating System: Windows 2003 SP1
  Hardware: BladeCenter HS20
  Software Installed: Tivoli Provisioning Manager, WebSphere Application Server 5.1, DB2 8.1, ITDS5.2, Cygwin 1.5.16-1, IBM Director Agent

The managed servers have the following characteristics:

- msnwin1.itso.ibm.com
  IP: 10.1.1.184
  Hardware: BladeCenter HS20
  Operating System: Windows 2003 SP1
  Software Installed: IBM Director Agent, EWLM Managed Servers, WebSphere Application Server 5.1.1, IHS 1.3.28, Cygwin 1.5.16-1
- msnwin2.itso.ibm.com
  IP: 10.1.1.185
  Hardware: BladeCenter HS20
  Operating System: Windows 2003 SP1
  Software Installed: Director Agents, EWLM Managed Servers, WebSphere Application Server 5.1.1, IHS 1.3.28

7.1.2 The installation flow

In this project, the installation steps were as follows:

1. Installation of the Virtualization Engine console component.
   a. Installation of the Tivoli Directory Server prerequisite; Tivoli Directory Server installs the LDAP server with a DB2 DB.
   b. Installation of the Virtualization Engine console using the Virtualization Engine console wizard, that installs and configures a WebSphere Application server, the LDAP, the LTPA, the ISC and the Virtualization Engine console bridge to the IBM Director.
2. Installation of the IBM Director component, using the Virtualization Engine installation wizard, that installs a database and the IBM Director Agent.
3. Installation of Enterprise Workload Manager (EWLM), using the Virtualization Engine installation wizard, with the following sub tasks:
   a. EWLM Domain Manager installation and configuration
   b. EWLM Managed Server installation and configuration
4. Installation of Tivoli Provisioning Manager with the following sub tasks:
   a. install the prereqs; Cygwin, Tivoli Directory Server, DB2 and WebSphere Application server
   b. post installation tasks, which is mainly to populate the data center model (DCM)

Our installation process begins by completing the Virtualization Engine installation worksheet. This worksheet contains information taken from the Tivoli Directory Server installation as well as additional information required by the Virtualization Engine console installation wizard. The full worksheet is available in the Information Center:

http://publib.boulder.ibm.com/infocenter/eserver/v1r1/en_US/info/esmcinfo/eicacinstallvec.htm

The Table 7-1 is a copy of the Virtualization Engine installation worksheet which we used for our installation.

Table 7-1 Virtualization Engine installation worksheet

Virtualization Engine Installation Worksheet	
LDAP server connection information	
Installation screen label	**Answer**
LDAP server host	winve21.itso.ibm.com
LDAP server port	389 (default)
Domain name that is used for directory updates (Administrator DN)	cn=root
Password associated with domain name that is used for directory updates	password
LDAP security settings	
Installation screen label	**Answer**
DN of user and group tree (Suffix DN)	o=ibm, c=us
user naming attribute	uid
Relative location of users	cn=users
User object class	inetOrgPerson
Group naming attribute	cn
Relative location of groups	cn=groups
Group object class	groupOfUniqueNames
Group membership attribute	uniqueMember
User bind information for LDAP server	
Installation screen label	**Answer**
Domain name for user and group lookup	uid=iscadmin, cn=users, o=ibm, c=us
Password associated with DN for user and group lookup	password
LTPA security settings	
Installation screen label	**Answer**
LTPA password	passowrd
LTPA time-out	120
SSO domain name	itso.ibm.com
TCP connection information for the ISC	
Installation screen label	**Answer**
ISC host name	vewin21.itso.ibm.com
Starting port	8421

Virtualization Engine Installation Worksheet	
Integrated Solution Console - LDAP security	
Installation screen label	**Answer**
Portal Administrator DN	uid=iscadmin, cn=users, o=ibm, c=us
short user name	iscadmin
Portal administrator password	password
Portal administrator group domain name	cn=iscadmins, cn=group, o=ibm, c=us
Short group name	iscadmins

The Virtualization Engine installation can be broken into four stages. The first three involve using the Virtualization Engine installation wizard to install individual system services components. The last involves installing Tivoli Provisioning Manager because it uses its own installation wizard and process. The stages are as follows:

- The first stage is described in 7.2, “Installing the Virtualization Engine console component” on page 168. During this stage, you install the Virtualization Engine Integrated Solutions Console, the Virtualization Engine console bridge to IBM Director, and the Virtualization Engine base support files.
- The second stage is described in 7.3, “Installing the IBM Director components” on page 183.
- The third stage is described in 7.4, “Installing Enterprise Workload Manager” on page 186 and involves the installation process for EWLM.
- The final stage is described in 7.5, “Installing Tivoli Provisioning Manager” on page 192. and steps through the unique installation process for Tivoli Provisioning Manager

Technically, the first three stages could occur at the same time if you plan on installing each component on one system. In our environment we are installing the Virtualization Engine console and bridge to IBM Director, EWLM, and Tivoli Provisioning Manager on separate systems.

Tip: While it is possible to install the Virtualization Engine console and bridge to IBM Director and EWLM on a single system, it is a good idea to assess system performance and determine your specific fault tolerance requirements when planning your Virtualization Engine infrastructure.

7.1.3 The pre-installation steps

Before starting the installation, you need to get the code and create user IDs with the right authorizations.

Preparing the installation files

You can obtain the Virtualization Engine CD-ROM package or system service component image files from IBM. In our project, we use downloaded image files. Each image file represents a particular Virtualization Engine system service component. We stage these image files on our Virtualization Engine console management server by extracting them into the C:\Install_Files\VEC directory as shown in Figure 7-3.

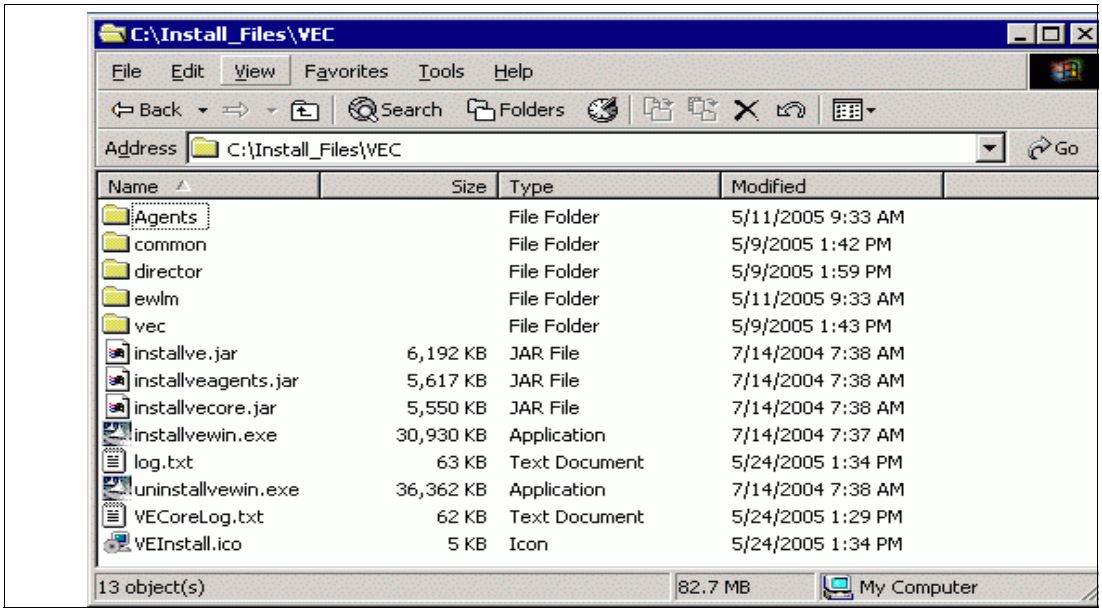
C:\Install_Files\VEC
File Edit View Favorites Tools Help
Back Search Folders
Address C:\Install_Files\VEC
Go
Name Size Type Modified
Agents File Folder 5/11/2005 9:33 AM
common File Folder 5/9/2005 1:42 PM
director File Folder 5/9/2005 1:59 PM
ewlm File Folder 5/11/2005 9:33 AM
vec File Folder 5/9/2005 1:43 PM
installve.jar 6,192 KB JAR File 7/14/2004 7:38 AM
installveagents.jar 5,617 KB JAR File 7/14/2004 7:38 AM
installvecore.jar 5,550 KB JAR File 7/14/2004 7:38 AM
installvewin.exe 30,930 KB Application 7/14/2004 7:37 AM
log.txt 63 KB Text Document 5/24/2005 1:34 PM
uninstallvewin.exe 36,362 KB Application 7/14/2004 7:38 AM
VECoreLog.txt 62 KB Text Document 5/24/2005 1:29 PM
VEInstall.ico 5 KB Icon 5/24/2005 1:34 PM
13 object(s) 82.7 MB My Computer

Figure 7-3 Virtualization Engine console staging area on vewin21.itso.ibm.com

Tip: You can stage the installation files anywhere within your network. Be sure you use physical drive mappings for this staging area if you plan on installing over the LAN. Errors will occur and your installation will fail if you try installing over the LAN using just a UNC path.

With the Virtualization Engine images staged in this format each Virtualization Engine component appears as an option when the Virtualization Engine installation wizard is launched, as shown in Figure 7-4.

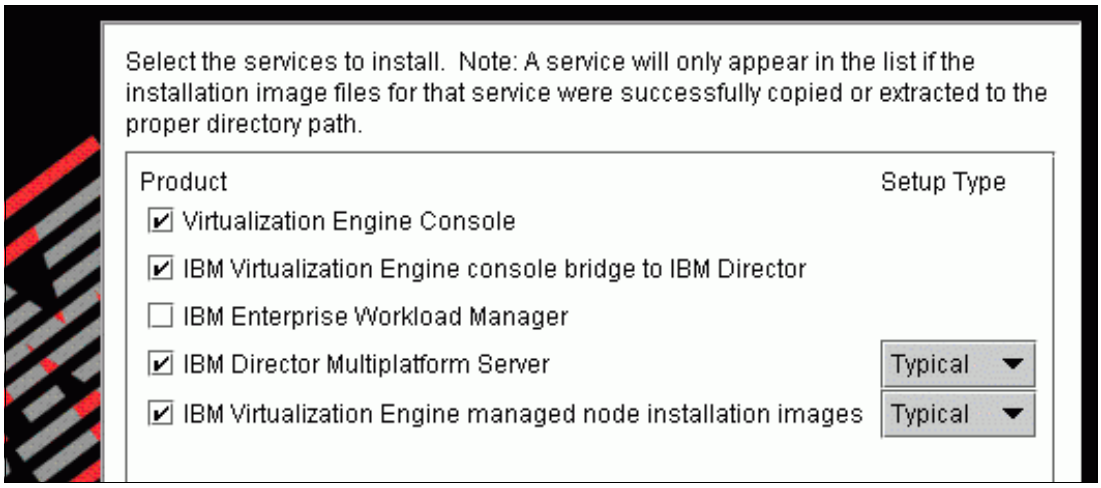
Select the services to install. Note: A service will only appear in the list if the installation image files for that service were successfully copied or extracted to the proper directory path.
Product
Setup Type
Virtualization Engine Console
IBM Virtualization Engine console bridge to IBM Director
IBM Enterprise Workload Manager
IBM Director Multiplatform Server
Typical
IBM Virtualization Engine managed node installation images
Typical

Figure 7-4 Virtualization Engine installation wizard

Note: In our environment, we installed each Virtualization Engine system service component in its default installation directory as follows:

- Virtualization Engine console, IBM Director, EWLM: C:\Program Files\IBM\VE
- Virtualization Engine console to IBM Director Bridge: C:\VECDir
- Tivoli Directory Server: C:\Program Files\IBM\LDAP
- Tivoli Provisioning Manager: C:\IBM\VE

Creating User Accounts

The Virtualization Engine installation requires several Windows user accounts. Some of these are automatically created by the installation wizard while others must be created manually. In the remaining sections of this chapter, we specify the accounts which need to be created manually.

7.2 Installing the Virtualization Engine console component

In this section we discuss the following:

- Installing the Virtualization Engine console Tivoli Directory Server 5.1 prerequisites.
- Installing the following components using the Virtualization Engine console installation wizard:
  - WebSphere 5.1 configuration settings
  - LDAP configuration settings
  - Lightweight third party authentication (LTPA) settings
  - Integrated Solutions Console configuration
  - Virtualization Engine console bridge to IBM Director configuration
  - Verifying the Virtualization Engine console installation
  - Logging into the ISC

7.2.1 Installing and configuring the Virtualization Engine console prerequisites

The Virtualization Engine console requires the following prerequisite: a server running IBM Tivoli Directory Server Version 5.1. We used the following default installation directory: C:\Program Files\IBM\LDAP.

This application is installed separately from the Virtualization Engine installation. As with the rest of the Virtualization Engine we installed Tivoli Directory Server using a downloaded image file. The following section describes the installation and configuration of Tivoli Directory Server 5.1 as it pertains to the Virtualization Engine.

The Tivoli Directory Server installation contains several steps. The following pages provide key screen-shots which capture critical user input fields. For detailed step-by-step instructions, refer to the instructions for installing Tivoli Directory Server at:

http://publib.boulder.ibm.com/infocenter/eserver/v1r1/en_US/info/esmcinfo/eicacldapinstallwindows.htm

Also see the *Tivoli Directory Server 5.1 Installation Guide* at:

http://publib.boulder.ibm.com/tividd/td/IBMDS/IDSapinst51/en_US/HTML/ldapinst.htm

Pre-installation considerations

Consider the following before you install:

1. Create a Windows user account with administrator privileges. This account is used by the database that stores the Tivoli Directory Server LDAP data. The user account name must be all lowercase letters. We created an account called itdsdb2.
2. Launch the Tivoli Directory Server setup.exe file. In our example this file is located in the C:\Install_Files\ITDS5.1\ids_ismp\ directory, as shown in Figure 7-5.

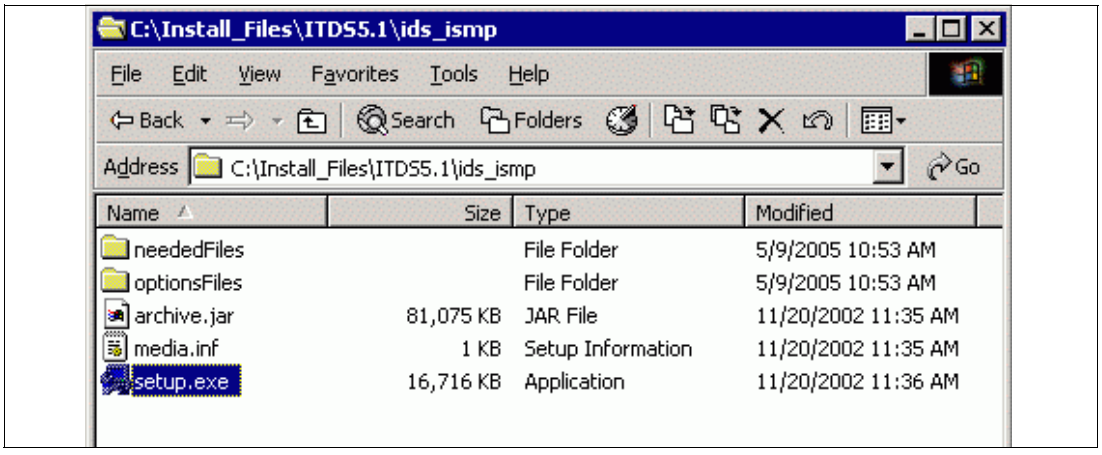
C:\Install_Files\ITDS5.1\ids_ismp
File Edit View Favorites Tools Help
Back Search Folders
Address C:\Install_Files\ITDS5.1\ids_ismp
Go
Name Size Type Modified
neededFiles File Folder 5/9/2005 10:53 AM
optionsFiles File Folder 5/9/2005 10:53 AM
archive.jar 81,075 KB JAR File 11/20/2002 11:35 AM
media.inf 1 KB Setup Information 11/20/2002 11:35 AM
setup.exe 16,716 KB Application 11/20/2002 11:36 AM

Figure 7-5 Tivoli Directory Server setup files

3. Figure 7-6 lists the components to be installed. Select the components that you want to install.

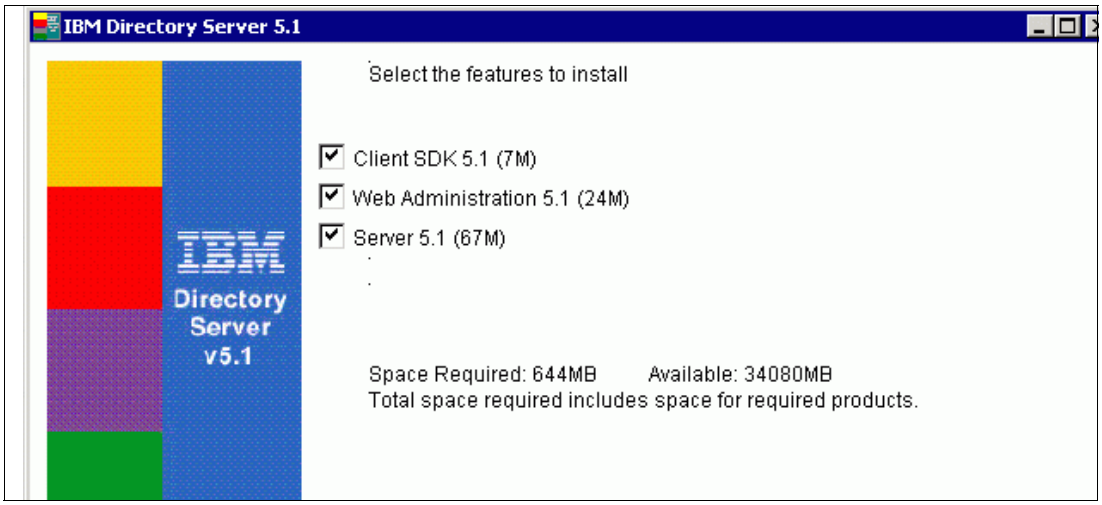

Figure 7-6 List of components

4. Tivoli Directory Server installs DB2 so it can store the LDAP directory information. Before the install begins we are prompted for a DB2 service account. The default account is db2admin. You can change this account name, or specify an account that already exists. The account requires administrator privileges. We chose the default db2admin account. Figure 7-7 shows the panel you need to complete.

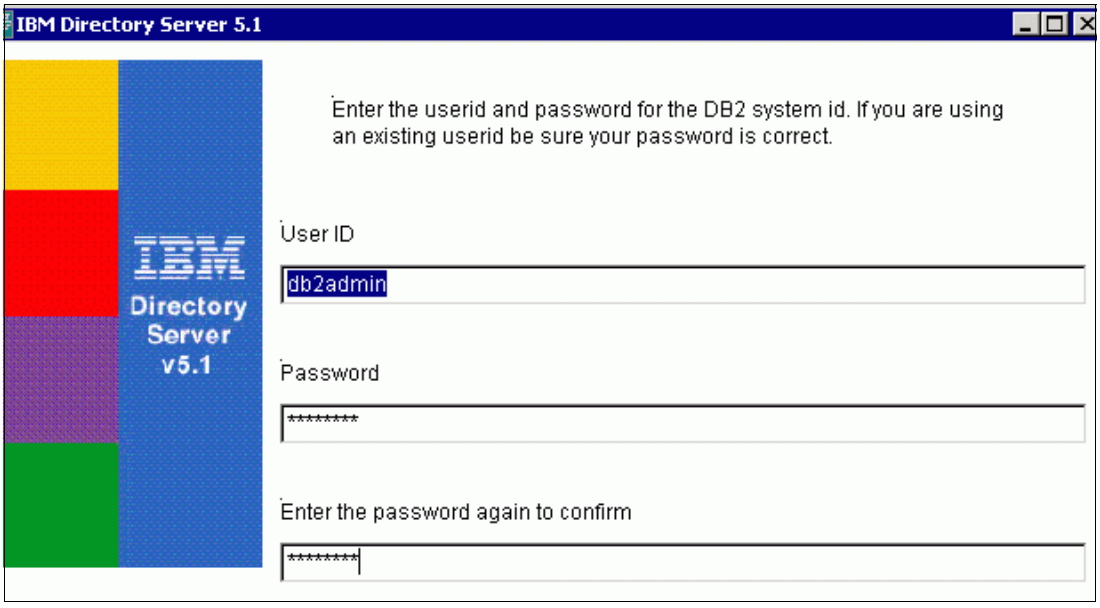

Figure 7-7 DB2 account

5. The Tivoli Directory Server installation installs DB2 for storing the LDAP database, the Global Security Toolkit (GSKit) for SSL and RSA support, and then the Web Administration interface which installs WebSphere 5.0 - Express.
6. When the installation completes, follow the prompts and restart your system.

Configuring Tivoli Directory Server

After Tivoli Directory Server is installed, the Tivoli Directory Server Administrator Distinguished Name (Admin DN) and the LDAP database are configured.

After the system restarts, the IBM Directory Configuration Tool is launched, as shown in Figure 7-8. If the tool does not start automatically or if you accidently close the tool, you can restart it by issuing the **`ldabxcfg`** command from the Windows command line.

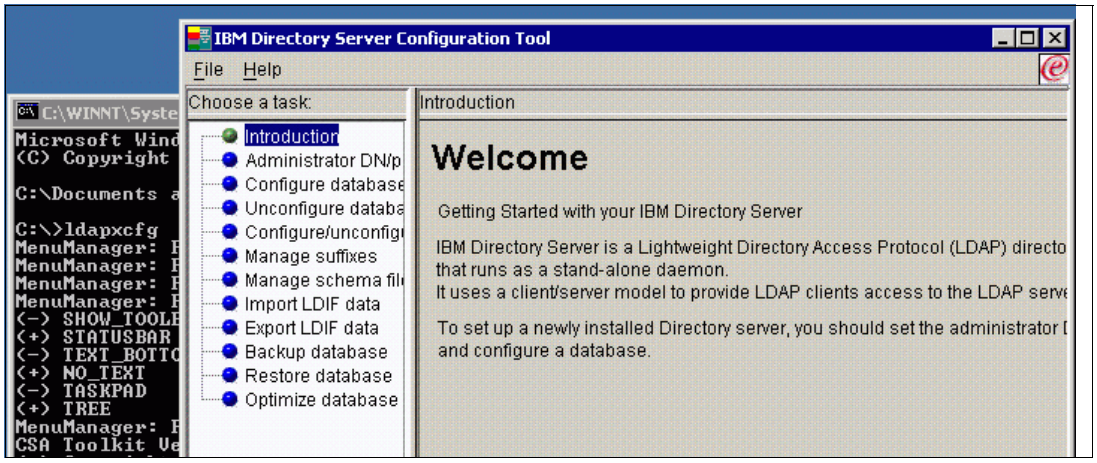

Figure 7-8 Tivoli Directory Server Configuration Tool

The Tivoli Directory Server installation involves setting some LDAP parameters for the environment. Refer to the LDAP primer located in the Information Center:

`http://publib.boulder.ibm.com/infocenter/eserver/v1r1/en_US/info/esmcinfo/eicacldapconcepts.htm`

Complete the following steps:

1. First, click **Administrator DN/password**. We use the default, `cn=root`. Enter a password, and then confirm the password, as shown in Figure 7-9 . Click **OK**.

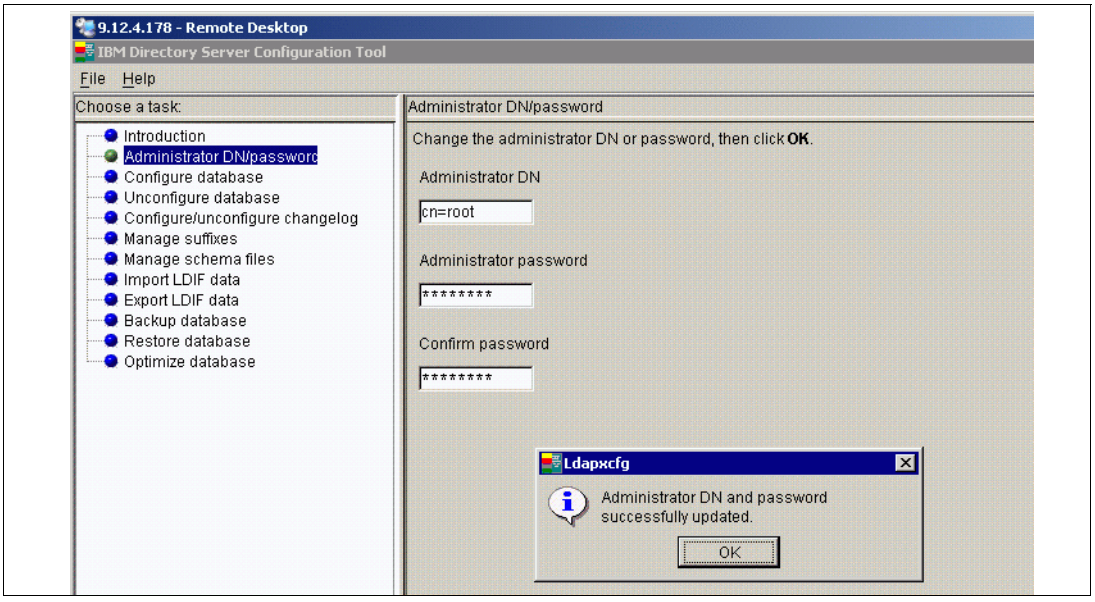

Figure 7-9 Adding the Administrator DN

2. Click **Configure database** and enter in the itdsdb2 user ID that you created earlier, as shown in Figure 7-10.

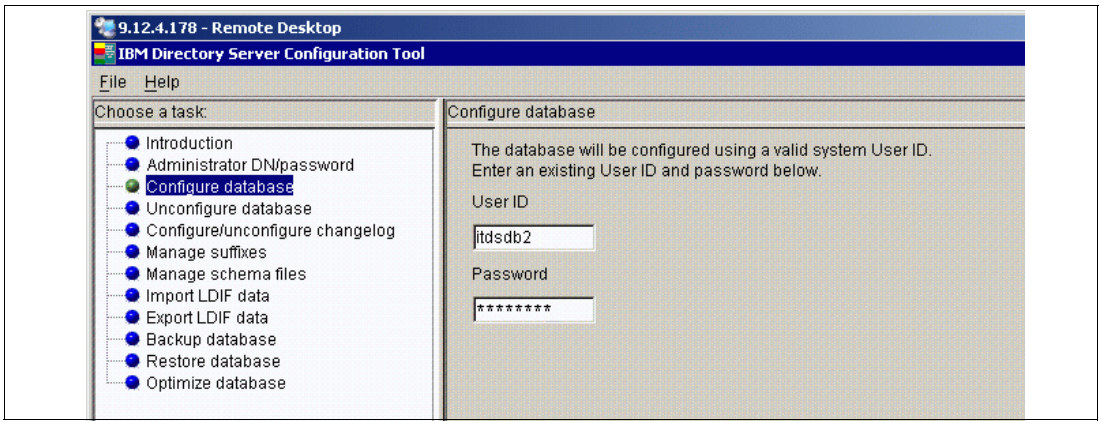

Figure 7-10 DB configuration panel

3. Specify the name of the LDAP database, as shown in Figure 7-11.

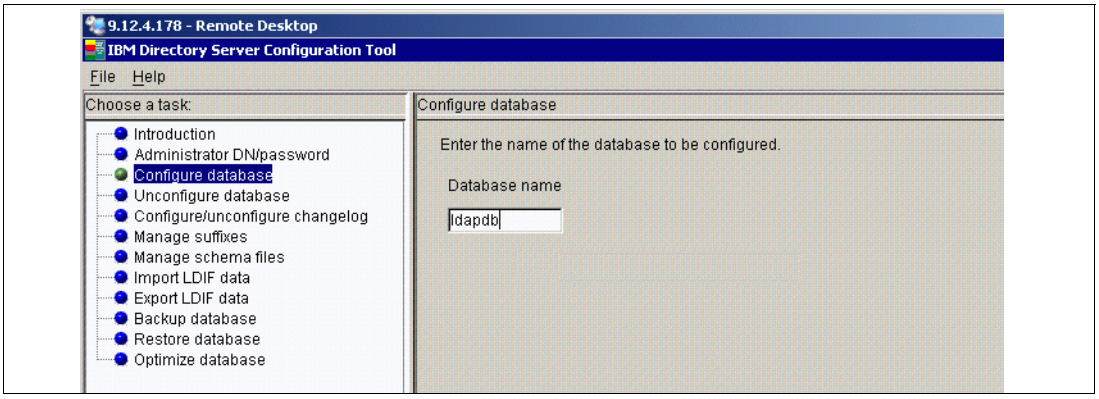

Figure 7-11 LDAP DB configuration panel

4. Select the database character set and location. We chose the default options. Click **Next** twice, and then **Finish**. The Tivoli Directory Server configuration tool builds the database. At the end, the information is summarized, as shown in Figure 7-12.

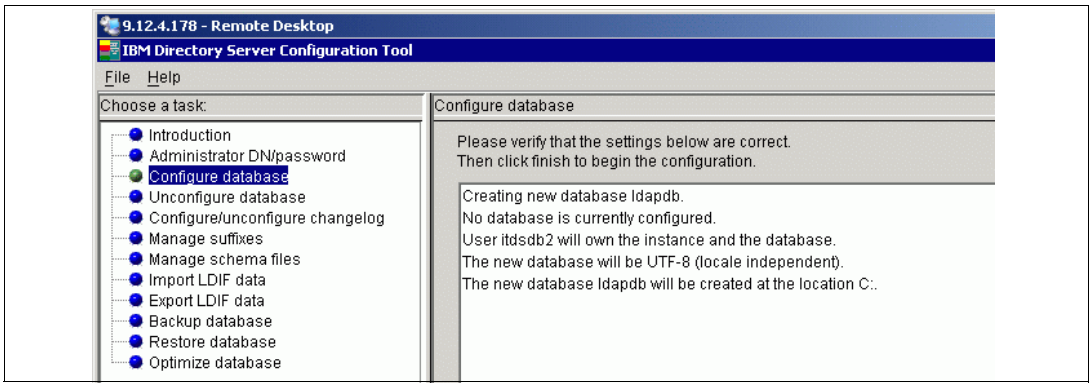

Figure 7-12 Installation summary

5. After the database is completed, click **Manage suffixes** as shown in Figure 7-13. Our DN suffix is o=ibm, c=us. When the suffix DN is added, close the Tivoli Directory Server configuration tool. The next step is starting the Tivoli Directory Server service.

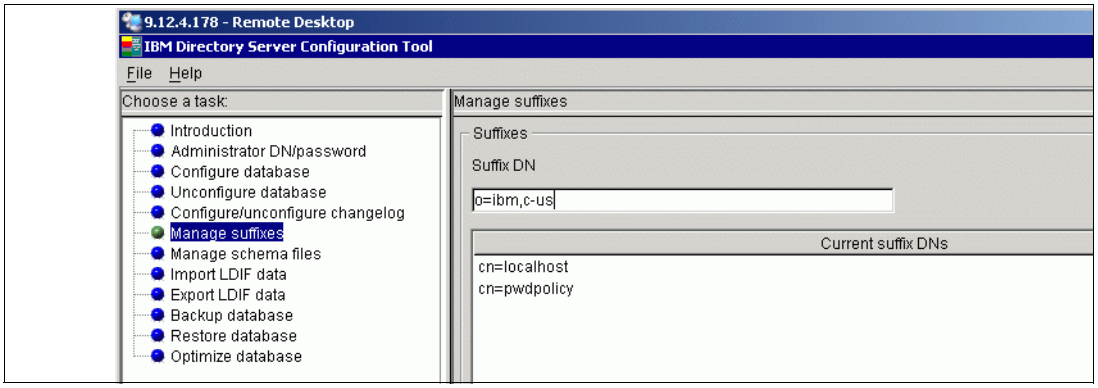

Figure 7-13 Managing suffixes

Starting the Tivoli Directory Server service

You can start the IBM Directory server via the Windows services control panel, as shown in Figure 7-14, or interactively via a Windows command line as shown in Example 7-1 on page 174.

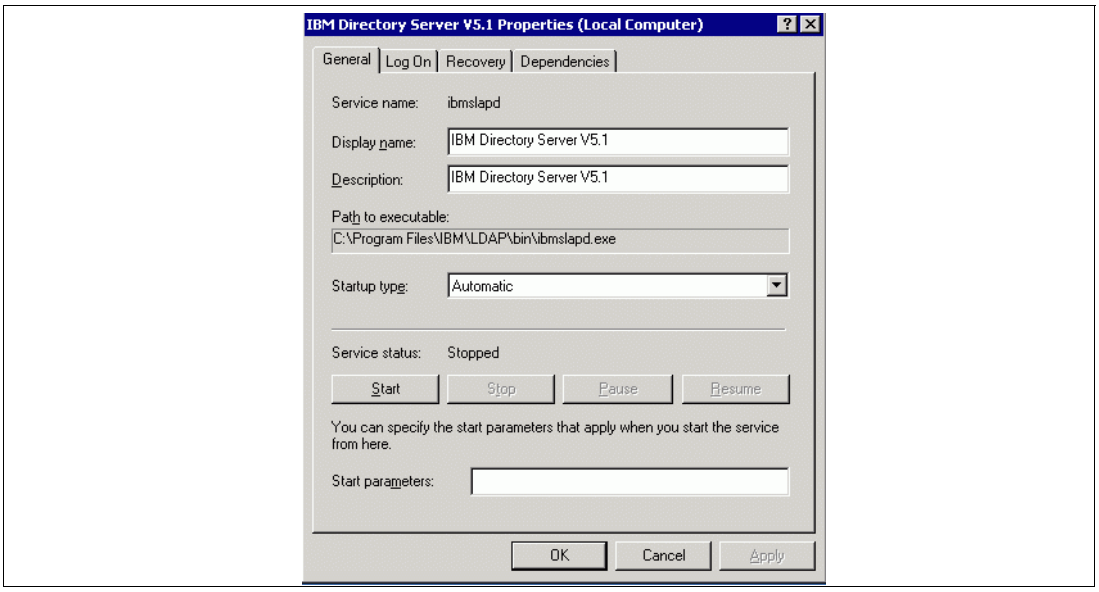

Figure 7-14 Starting Tivoli Directory Server via the Windows service control

Example 7-1 Starting Tivoli Directory Server via the Windows command line

```
C:\Program Files\IBM\LDAP\bin>ibmslapd
Server starting.
Plugin of type EXTENDEDOP is successfully loaded from libevent.dll.
Plugin of type EXTENDEDOP is successfully loaded from libtranext.dll.
Plugin of type EXTENDEDOP is successfully loaded from libldaprepl.dll.
Plugin of type PREOPERATION is successfully loaded from libDSP.dll.
Plugin of type EXTENDEDOP is successfully loaded from libevent.dll.
Plugin of type EXTENDEDOP is successfully loaded from libtranext.dll.
Plugin of type AUDIT is successfully loaded from C:/Program Files/IBM/LDAP/bin/l
ibldapaudit.dll.
Plugin of type EXTENDEDOP is successfully loaded from libevent.dll.
Plugin of type EXTENDEDOP is successfully loaded from libtranext.dll.
Plugin of type DATABASE is successfully loaded from C:/Program Files/IBM/LDAP/bi
n/libback-rdbm.dll.
Plugin of type REPLICATION is successfully loaded from C:/Program Files/IBM/LDAP
/bin/libldaprepl.dll.
Plugin of type EXTENDEDOP is successfully loaded from libevent.dll.
Plugin of type DATABASE is successfully loaded from C:/Program Files/IBM/LDAP/bi
n/libback-config.dll.
Plugin of type EXTENDEDOP is successfully loaded from libloga.dll.
Non-SSL port initialized to 389.
IBM Directory (SSL), Version 5.1        Server started.
```

Tip: Tivoli Directory Server needs its DB2 database online so it can start. If you configure Tivoli Directory Server to start automatically, after a reboot you might get an error stating that the IBM Directory Server V5.1 service hung on starting. This error occurs because Tivoli Directory Server issues a call to start its database instance and it takes longer than expected for the database to come online. However, when you look at the service in the service control manager, Tivoli Directory Server is in the running state. Ignore the message, or start the service manually.

For a production environment we recommend to create dependencies, using the Dependencies tag as shown in Figure 7-14 on page 173.

Now that Tivoli Directory Server is running, the next section describes how to create the LDAP user and group accounts for the Virtualization Engine console administrator. This is done using the Tivoli Directory Server Web administration tool. When these Virtualization Engine console administrator user and group accounts are created, the Tivoli Directory Server configuration will be complete and the Virtualization Engine console authentication will be established.

Creating LDAP user and group accounts for the Virtualization Engine Console Administrator

To create the LDAP user and group accounts, do the following:

1. Start the Tivoli Directory Server Web administration WebSphere instance, using the **startserver** script as shown in Example 7-2.

Example 7-2 Starting the Tivoli Directory Server Web administration tool

```
C:\Program Files\IBM\LDAP\appsrv\bin>startserver server1
ADMU0116I: Tool information is being logged in file C:\Program
           Files\IBM\LDAP\appsrv\logs\server1\startServer.log
ADMU3100I: Reading configuration for server: server1
ADMU3200I: Server launched. Waiting for initialization status.
ADMU3000I: Server server1 open for e-business; process id is 3780

C:\Program Files\IBM\LDAP\appsrv\bin>
```

2. Log in to the Tivoli Directory Server Web Administration Console using the following link, as shown in Figure 7-15:

   http://localhost:9080/IDSWebApp/IDSjsp/Login.jsp

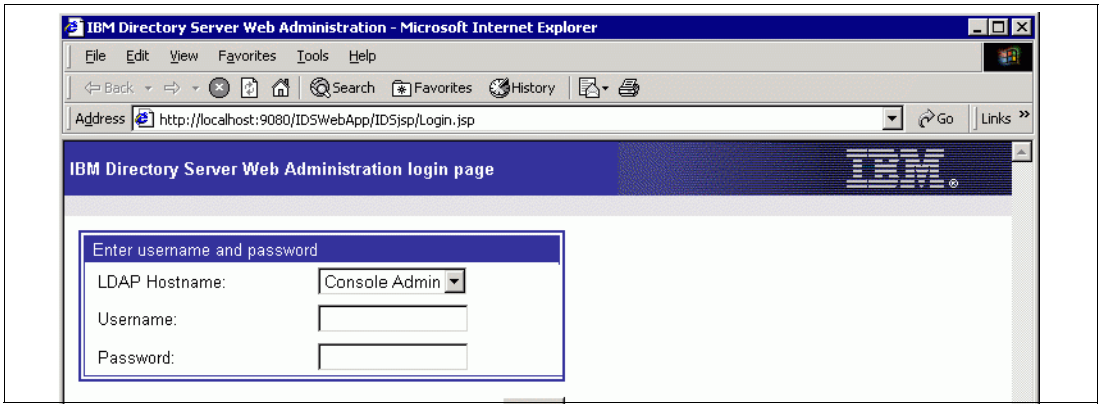

Figure 7-15 Log in to the Tivoli Directory Server Web administration console

Tip: The Tivoli Directory Server WebSphere instance is not configured as a Windows service. If your system is rebooted the Tivoli Directory Server WebSphere instance can be restarted manually. Use the startup procedure detailed in Example 7-2.

3. Log in to the Tivoli Directory Server Web Administration tool with the built in Console Admin user ID, superadmin. The password for this account is secret.
4. As shown in Figure 7-16 on page 176, click **Console Administration** → **Manage Console Servers**, and click **Add**. Enter the fully qualified host name of your Tivoli Directory Server server leaving the remaining fields at their default values, and click **OK**.

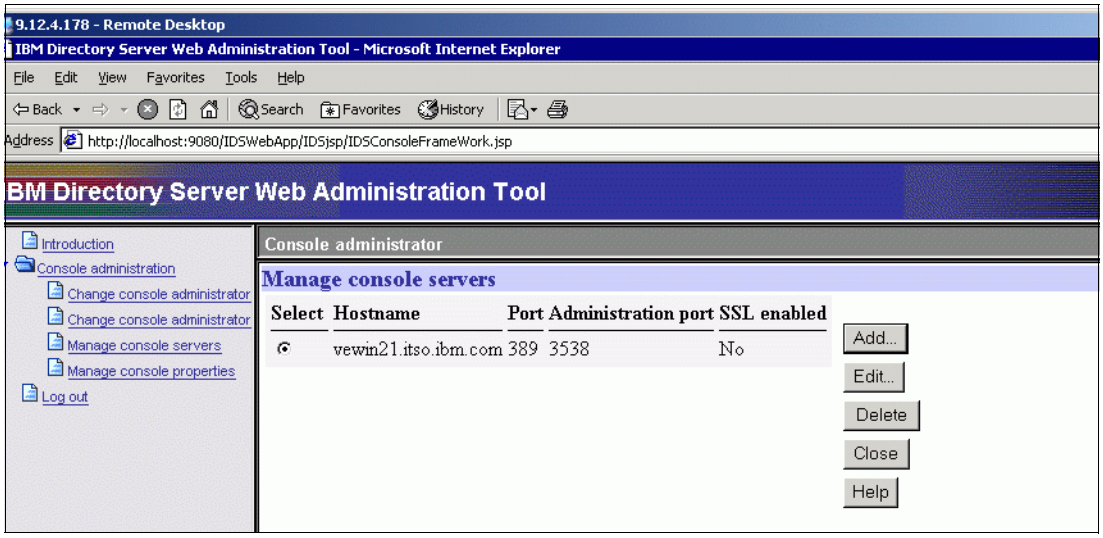

Figure 7-16 Managing console servers

5. Log out of the Tivoli Directory Server Web Administration tool. Then log back in, this time changing the LDAP hostname field to the server added in the last step . Use the Administrator DN, `cn=root`, and password that you created earlier.

The next steps involve setting up the LDAP directory and creating the Virtualization Engine console administrator account:

1. From the main Tivoli Directory Server Web Administration tool screen, click **Directory Management** → **Add and entry**.
2. In the Select Object Class frame, scroll down to organization and select **Next**.
3. Click **Next** and leave the auxiliary object classes blank. On the next screen, enter the Distinguished Name for your object class, as shown in Figure 7-17. The DN consists of a Relative DN (`o=ibm`) and a Parent DN (`c=us`). These values were specified earlier as the Suffix DN during the Tivoli Directory Server Configuration.

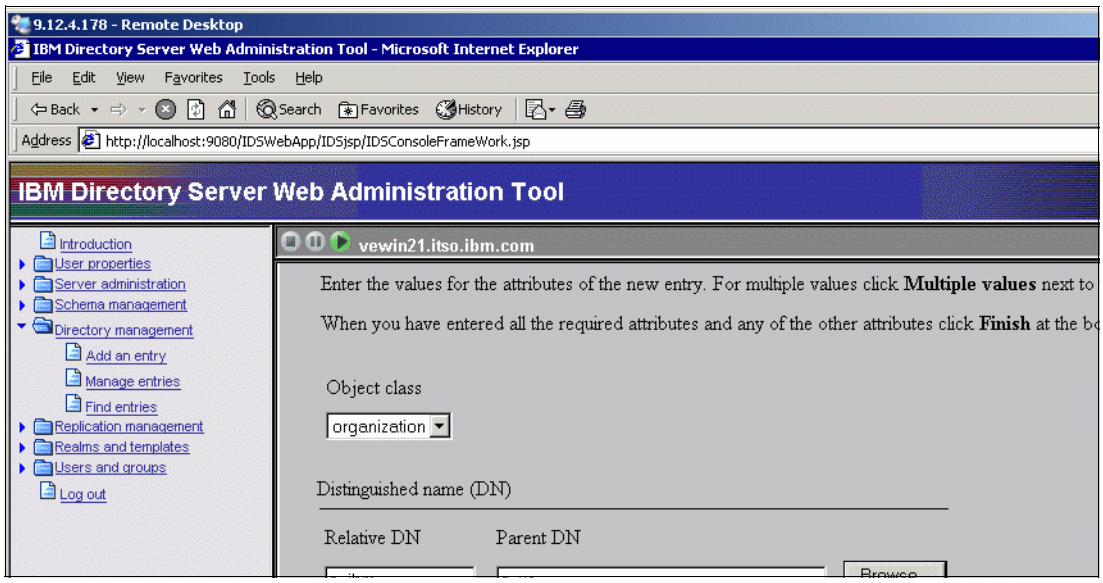
9.12.4.178 - Remote Desktop
IBM Directory Server Web Administration Tool - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites History
Address http://localhost:9080/IDSWebApp/IDSjsp/IDSConsoleFrameWork.jsp
IBM Directory Server Web Administration Tool
Introduction
User properties
Server administration
Schema management
Directory management
Add an entry
Manage entries
Find entries
Replication management
Realms and templates
Users and groups
Log out
vewin21.itso.ibm.com
Enter the values for the attributes of the new entry. For multiple values click Multiple values next to
When you have entered all the required attributes and any of the other attributes click Finish at the bo
Object class
organization
Distinguished name (DN)
Relative DN
Parent DN
Browse

Figure 7-17 Create the Virtualization Engine console administrator account step 1

4. Click **Finish** → **Directory Management** → **Manage Entries**.
5. Select the entry added in the previous step (`o=ibm, c=us`), and click **Add**. A panel similar to that shown in Figure 7-18 is then displayed.

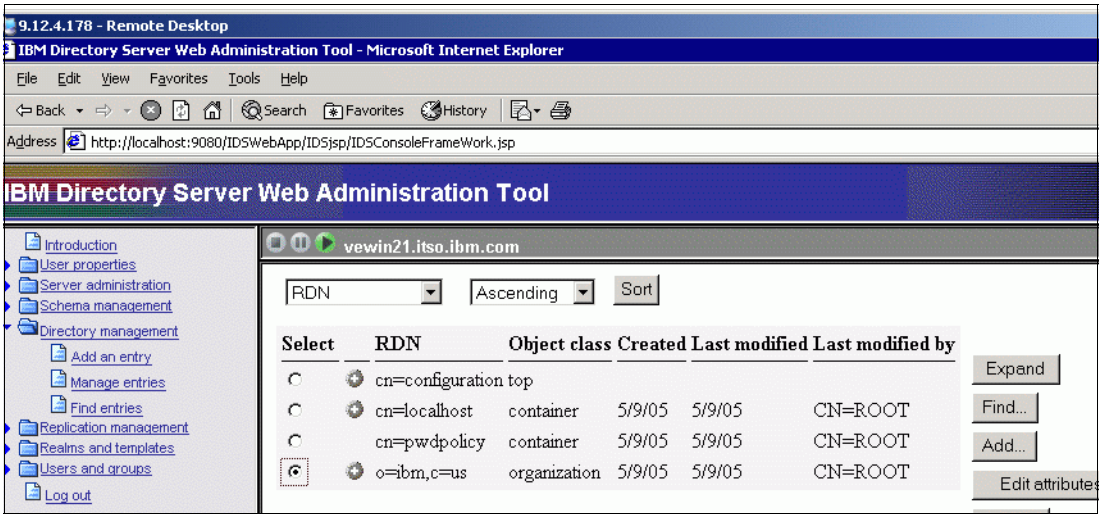

Figure 7-18 Create the Virtualization Engine console administrator account step 2

6. Choose **container** in the Select object class frame, and click **Next**. You need then to complete the panel as shown in Figure 7-19.
   - Enter `cn=users` in the Relative DN field. The Parent DN field is already populated with the suffix DN o=ibm,c=us.
   - Leave the Other attributes field blank and click **Finish**.

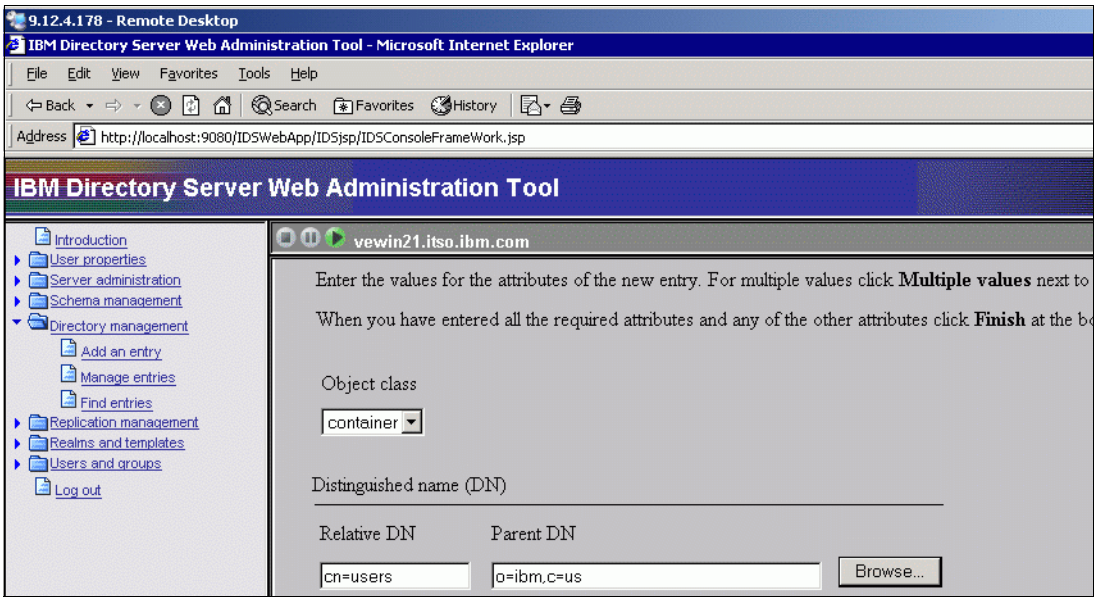

Figure 7-19 Adding cn=users container

7. Now repeat the previous step and create an additional container object class entry called cn=groups.

8. Create the Virtualization Engine console Administrator account called iscadmin, as shown in Figure 7-20 by doing the following:
   a. Click **Directory Management** → **Manage Entries**.
   b. Select the **o=bim,c=us** DN entry as in the previous steps, and click **Expand**.
   c. Click **cn=users**, and click **Add**.
   d. Choose **inetOrgPerson** in the Select object class field, and click **Next**.

      In the relative DN specify the user account for administering the Virtualization Engine Console, cn=iscadmin. The parent DN is already populated with cn=users,o=ibm,c=us.
   e. Enter `iscadmin` in both the cn and sn fields at the bottom under Required attributes.
   f. Click **Other attributes** (even though it does not look like a button or link). Scroll to the userPassword field and enter a password for the iscadmin account. Then click **Finish**.

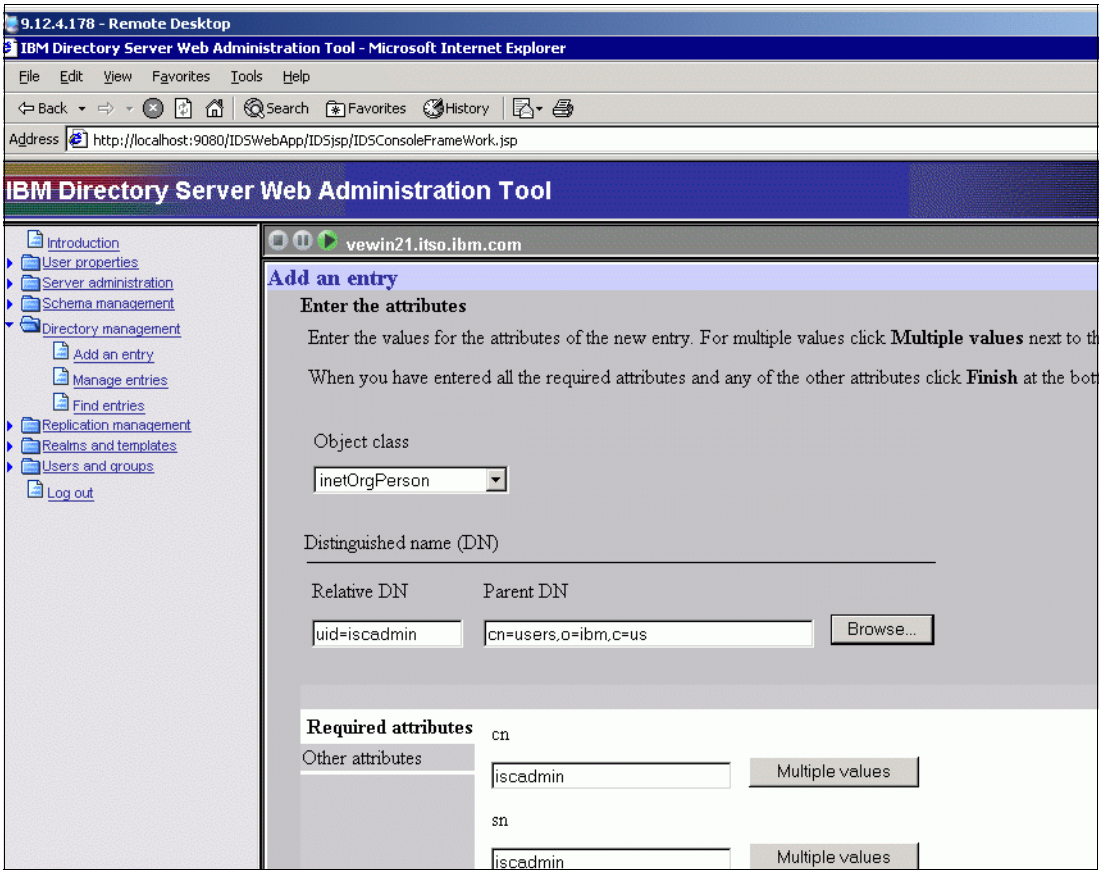

Figure 7-20 Adding the iscadmin user account

9. Create a group for the iscadmin account that you created in the previous step, as shown in Figure 7-21 on page 179 by doing the following:
   a. Click **Directory Management** → **Manage Entries**.
   b. Select the **o=bim,c=us** DN entry as in the previous steps, and click **Expand**.
   c. Select **cn=groups**, and click **Add**.
   d. Choose **groupOfUniqueName** in the Select object class field, and click **Next**.
   e. In the relative DN field, specify a group for administering the Virtualization Engine Console (cn=iscadmins). The parent DN is already populated with cn=groups,o=ibm,c=us.

f. Type `iscadmins` in the cn field under Required attributes. Then type `uid=iscadmin,cn=users,o=ibm,c=us` as the DN for the iscadmin user account in the uniqueMember field.

g. Click **Finish**.

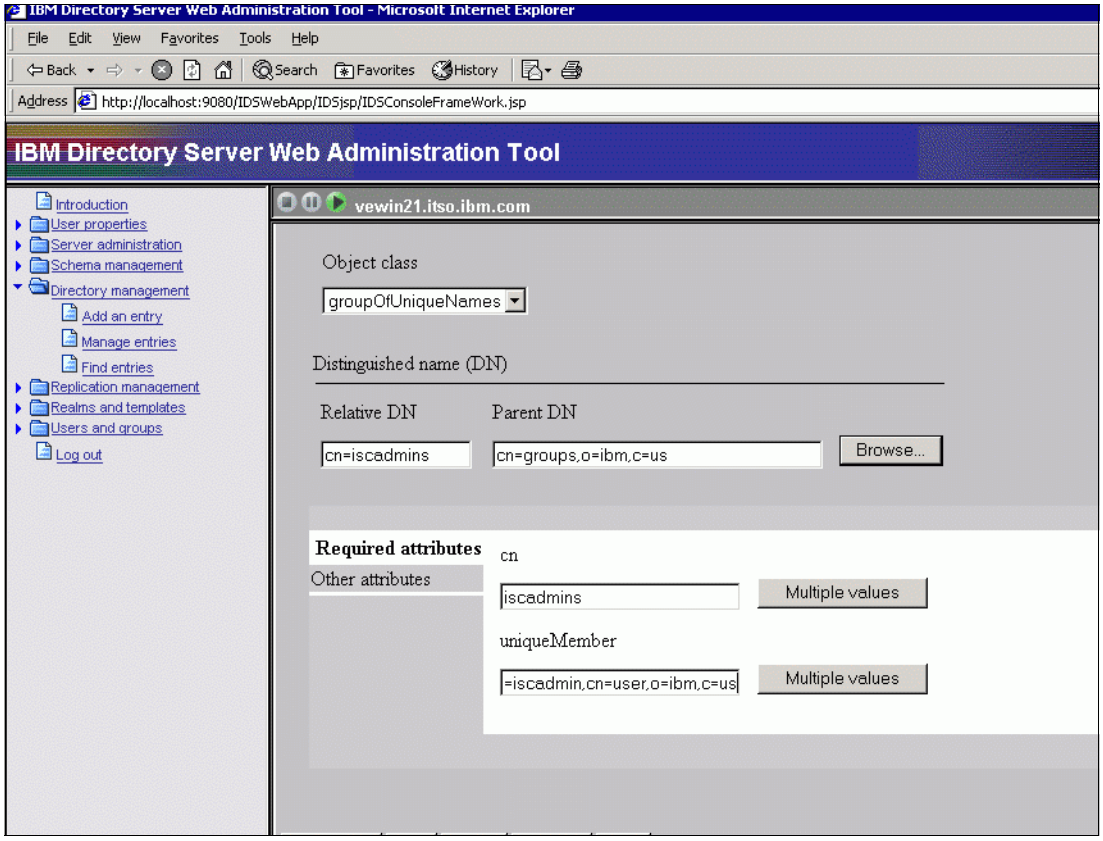

Figure 7-21 Creating the iscadmins group

The Tivoli Directory Server configuration is now complete. The system is ready for the Virtualization Engine console and Virtualization Engine system service components installation wizard.

7.2.2 Installing the Virtualization Engine console code using the wizard

In this section, we install the Virtualization Engine Integrated Systems Console, the Virtualization Engine bridge to IBM Director, and base Virtualization Engine support files (these are installed by default), as shown in Figure 7-22.

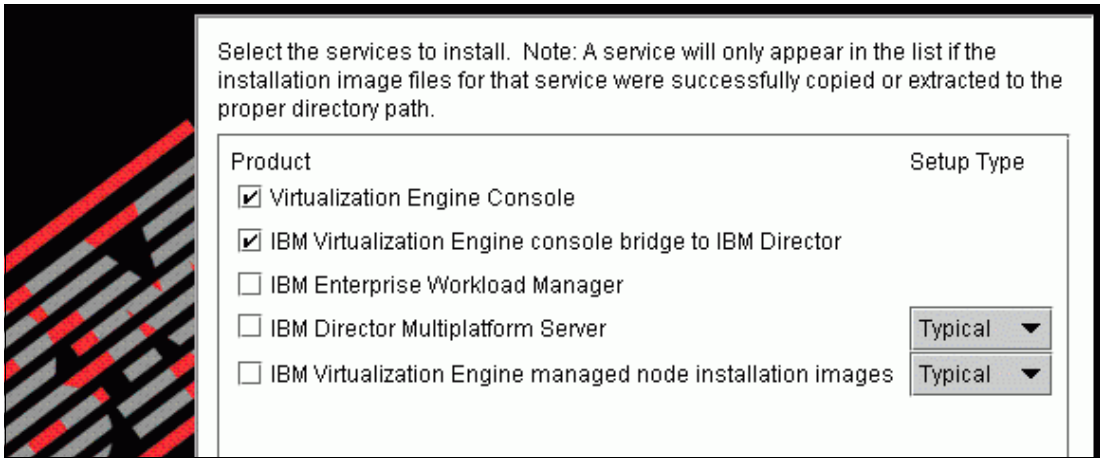
Select the services to install. Note: A service will only appear in the list if the installation image files for that service were successfully copied or extracted to the proper directory path.
Product
Setup Type
Virtualization Engine Console
IBM Virtualization Engine console bridge to IBM Director
IBM Enterprise Workload Manager
IBM Director Multiplatform Server
Typical
IBM Virtualization Engine managed node installation images
Typical

Figure 7-22 The list of services

Launch the Virtualization Engine installation wizard by executing the installvewin.exe file in your Virtualization Engine console staging directory. In our case this is C:\Install_Files\VEC \installvewin.exe. Accept all agreements and choose the appropriate language for the wizard. The panels provided by the wizard guide you through the installation of the Virtualization Engine console.

The installation is broken into several steps:

1. WebSphere 5.1 configuration settings

   Specify your WebSphere Application Server 5.1 server information and click **Next**.

2. LDAP configuration settings

   Specify the LDAP server information obtained during the Tivoli Directory Server installation. Be sure to use `cn=root` as the DN used for directory updates.

3. Lightweight third-party authentication (LTPA) settings

   Create a password for building the LTPA security keys. Add you domain suffix in the SSO domain name field.

4. Integrated Solutions Console configuration

   Specify connection and LDAP information for the ISC.

   Specify a location for the LTPA key file. The default is C:\Program Files\IBM\VE\VEConsole\LTPAKeys.dat. Make a backup of this key file and remember its location. This is used by the Virtualization Engine console to configure the WebSphere communication.

5. Virtualization Engine console bridge to IBM Director configuration

   The Virtualization Engine console Bridge to IBM Director provides Virtualization Engine console connectivity to the IBM Director data. This allows a systems administrator to log into the ISC and view system information obtained from Director. The bridge can be installed on any system in your environment. The default installation directory is C:\VECDir.

6. Complete the screen even though you have not yet installed the IBM Director server. You will create the user ID later. Do not test connectivity at this time, because it will fail.

At the end of the installation, an installation summary appears as in Figure 7-23.

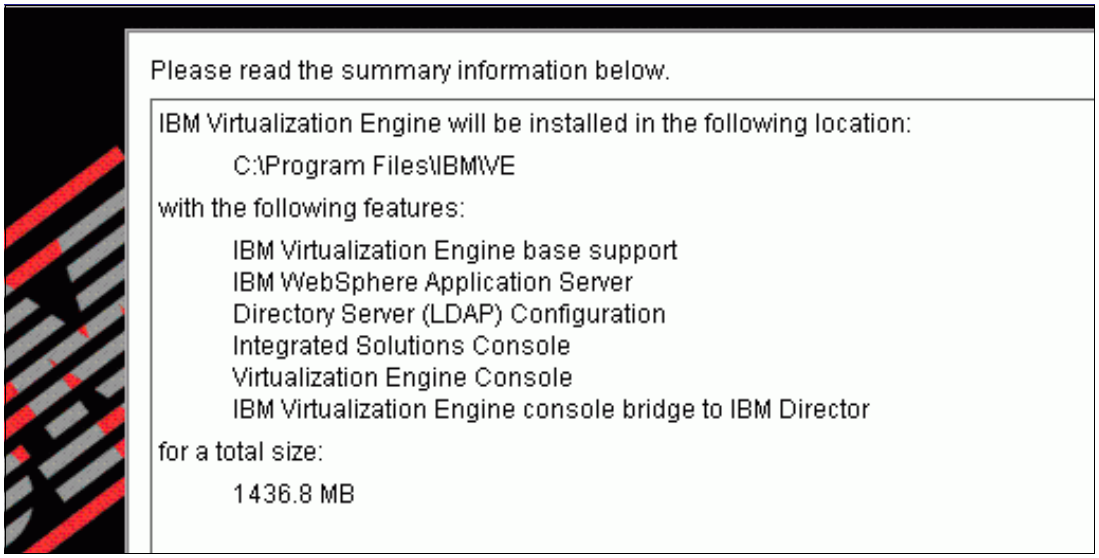
Please read the summary information below.
IBM Virtualization Engine will be installed in the following location:
C:\Program Files\IBM\VE
with the following features:
IBM Virtualization Engine base support
IBM WebSphere Application Server
Directory Server (LDAP) Configuration
Integrated Solutions Console
Virtualization Engine Console
IBM Virtualization Engine console bridge to IBM Director
for a total size:
1436.8 MB

Figure 7-23 Summary panel

Verifying the Virtualization Engine console installation

You can test the Virtualization Engine console installation by logging into the Integrated Solutions Console (ISC) as follows:

1. Open your Web browser and connect to the ISC at:

   `http://vewin21.itso.ibm.com:8421/ibm/defaultconsole/!ut/p/.scr/Login`

2. A panel as shown in Figure 7-24 must be displayed.

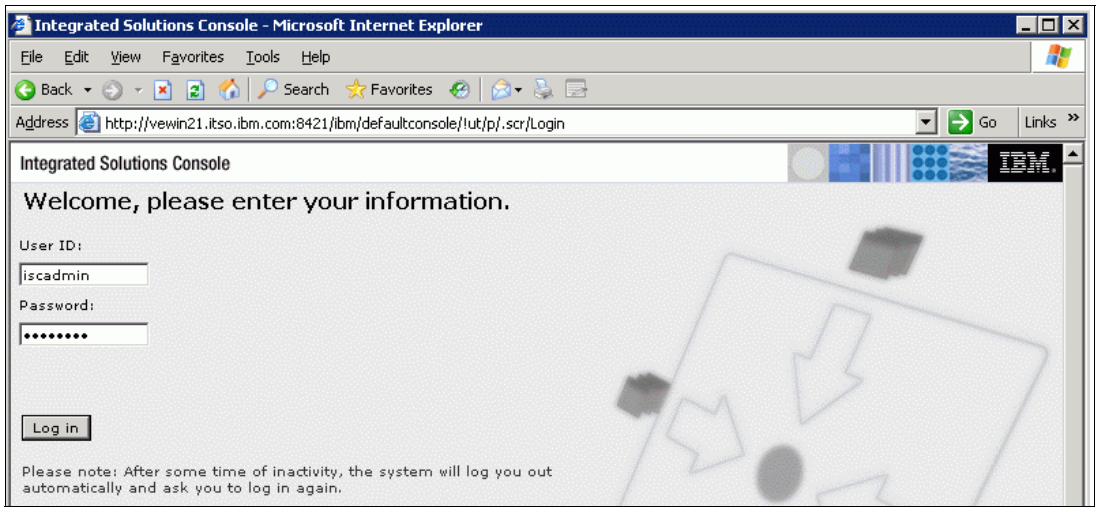
Integrated Solutions Console - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites
Address http://vewin21.itso.ibm.com:8421/ibm/defaultconsole/!ut/p/.scr/Login
Go Links
Integrated Solutions Console
IBM
Welcome, please enter your information.
User ID:
iscadmin
Password:
Log in
Please note: After some time of inactivity, the system will log you out automatically and ask you to log in again.

Figure 7-24 Logging to the ISC

3. After logging in, the windows shown in Figure 7-25 is displayed.

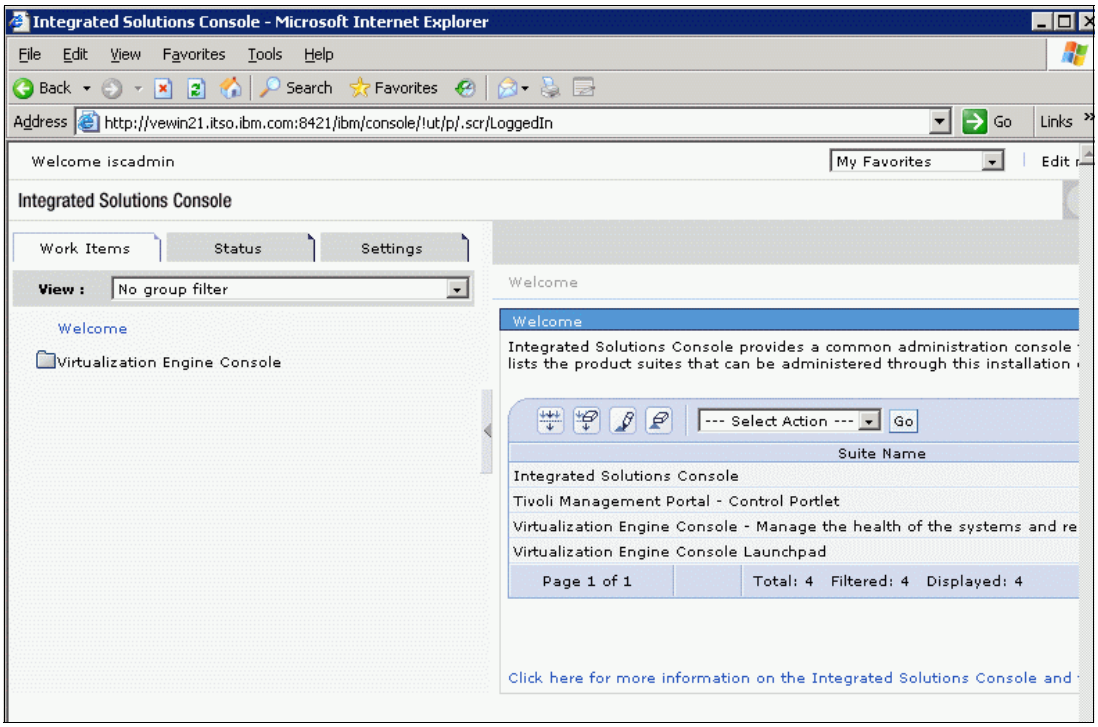
Integrated Solutions Console - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites
Address http://vewin21.itso.ibm.com:8421/ibm/console/!ut/p/.scr/LoggedIn
Go Links
Welcome iscadmin
My Favorites
Integrated Solutions Console
Work Items
Status
Settings
View : No group filter
Welcome
Virtualization Engine Console
Welcome
Integrated Solutions Console provides a common administration console
lists the product suites that can be administered through this installation
--- Select Action ---
Go
Suite Name
Integrated Solutions Console
Tivoli Management Portal - Control Portlet
Virtualization Engine Console - Manage the health of the systems and re
Virtualization Engine Console Launchpad
Page 1 of 1
Total: 4 Filtered: 4 Displayed: 4
Click here for more information on the Integrated Solutions Console and

Figure 7-25 The Virtualization Engine console primary panel

Remarks

The components of the Virtualization Engine console create several instances of WebSphere Application server. You can display the WebSphere version installed on your system by running the **versioninfo** (for WebSphere Application Server) or **wpversioninfo** (for WebSphere Application Server Portal) command in the installation directory \bin folder.

The Virtualization Engine console installation wizard also creates three Windows services as shown in Example 7-3.

Example 7-3 Windows services created by the Virtualization Engine console installation

```
IBM VE Console Help
IBM WebSphere Application Server V5 - WebSphere for VE Console
IBM WebSphere Application Server V5 - VECDir
```

These services have implicit dependencies upon each other that should be taken into consideration when planning and implementing your Virtualization Engine infrastructure.

For example, during the Virtualization Engine installation, two WebSphere services are created and both authenticate via Tivoli Directory Server. Tivoli Directory Server has a tendency to hang during automatic startup, causing the two WebSphere services to fail if they are also configured for automatic startup. This failure takes several minutes before timing out, during which time your system might be inaccessible.

Tip: In our case, we left the WebSphere Application Server services in a manual startup mode. After a reboot, log in to the server, confirm that Tivoli Directory Server is started, then manually start these services. Alternatively consider running Tivoli Directory Server on a separate system.

7.3 Installing the IBM Director components

This section explains how to install the IBM Director components through the Virtualization Engine wizard.

7.3.1 Choosing the components

To choose the components, follow these steps:

1. Start the wizard by executing the “installvewin.exe file in your Virtualization Engine console staging directory. In our case this is C:\Install_Files\VEC \installvewin.exe. Accept all agreements and choose the appropriate language for the wizard, as shown in Figure 7-26.

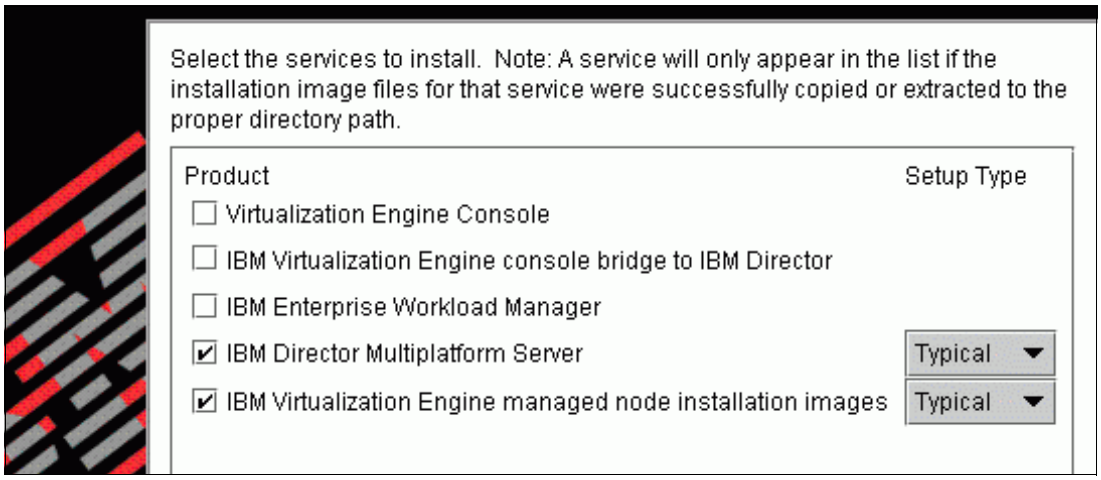
Select the services to install. Note: A service will only appear in the list if the installation image files for that service were successfully copied or extracted to the proper directory path.
Product
Setup Type
Virtualization Engine Console
IBM Virtualization Engine console bridge to IBM Director
IBM Enterprise Workload Manager
IBM Director Multiplatform Server
Typical
IBM Virtualization Engine managed node installation images
Typical

Figure 7-26 List of components to install using the Virtualization Engine wizard

2. Click **Next** twice and select the IBM Director features that you want to install as shown in Figure 7-27 on page 184. We choose the default.

Important: You cannot re-run the Virtualization Engine installation wizard and add additional IBM Director features. Be sure you choose all the features you need the first time.

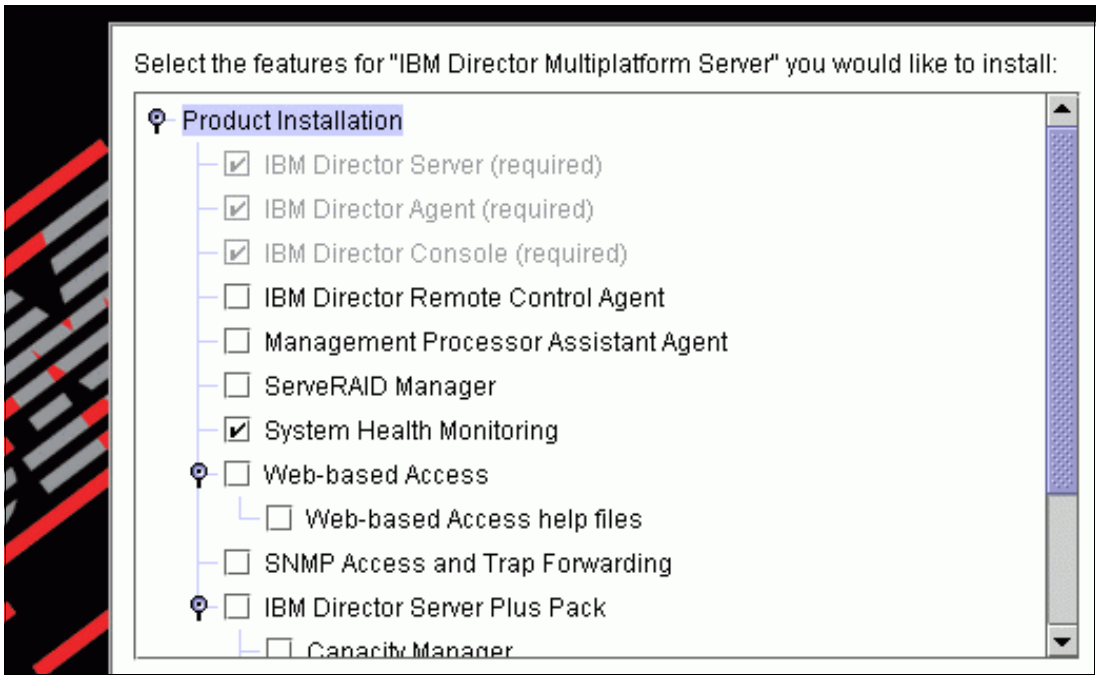
Select the features for "IBM Director Multiplatform Server" you would like to install:
Product Installation
IBM Director Server (required)
IBM Director Agent (required)
IBM Director Console (required)
IBM Director Remote Control Agent
Management Processor Assistant Agent
ServeRAID Manager
System Health Monitoring
Web-based Access
Web-based Access help files
SNMP Access and Trap Forwarding
IBM Director Server Plus Pack
Capacity Manager

Figure 7-27 List of subcomponents

3. Click **Next** and enter the IBM Director service account user ID that you created earlier.

   During the Virtualization Engine console bridge installation, you were prompted for an IBM Director Service Account. You need to configure the account that you specified at the time (User ID: ibmdirecto Password: password), as shown in Figure 7-28.

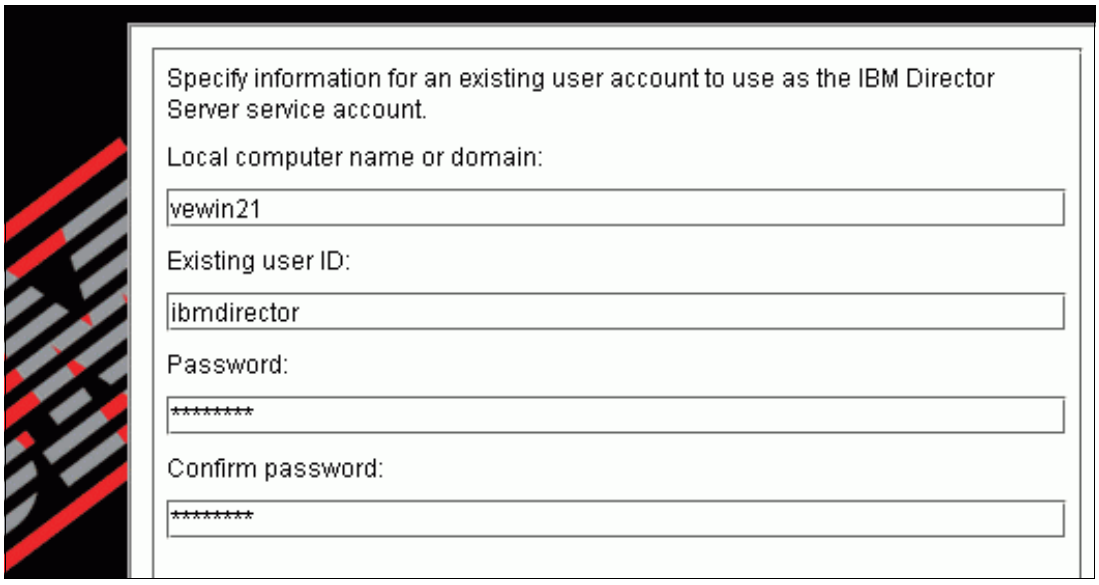

Figure 7-28 Create account userid

- Click **Next** and specify the type of encryption used for transmitting data between the Director server and the Director agents.
- The final screens summarize the installation information just provided. Then the installation begins. Upon completion reboot your server.

7.3.2 Creating the IBM Director Database

Director stores the server information it collects in a relational database. Supported options include IBM DB2, Microsoft Jet 4.0, Microsoft SQL Server 2000, Microsoft Data Engine 2000, Oracle Server, or PostgreSQL. We used IBM DB2 which is already running on our management server as a result of the Tivoli Directory Server installation.

Important: IBM Director requires DB2 UDB v8.1 Fix Pack 5. We had to upgrade the DB2 installation created by Tivoli Directory Server 5.1 because it was at the base DB2 8.1 level.

To create the IBM Director Database, follow these steps:

1. Launch the database configuration wizard by running the **`cfgdb`** command from the Windows command prompt. Choose the type of database and click **Next** as shown in Figure 7-29.

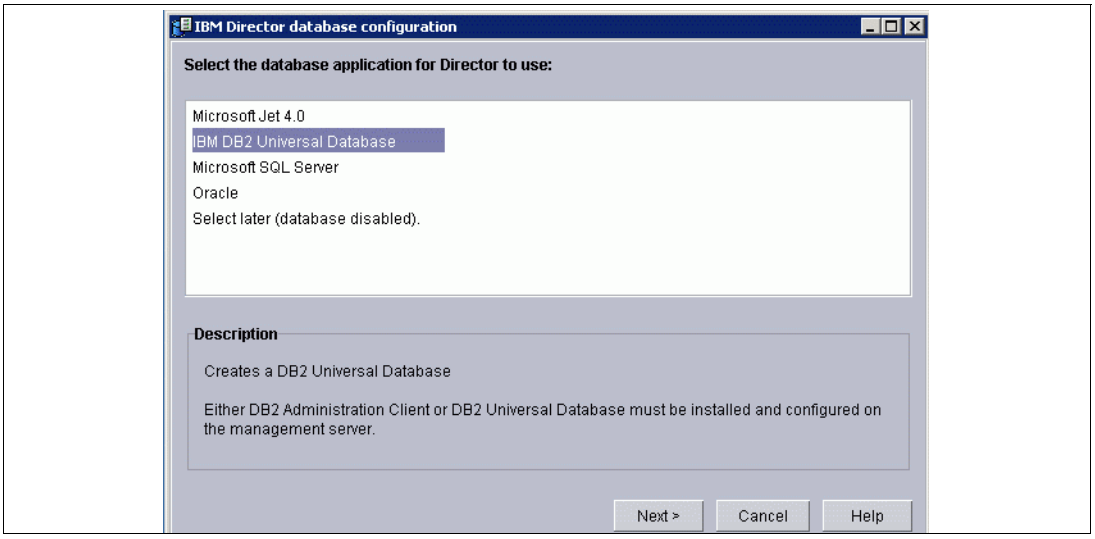
IBM Director database configuration
Select the database application for Director to use:
Microsoft Jet 4.0
IBM DB2 Universal Database
Microsoft SQL Server
Oracle
Select later (database disabled).
Description
Creates a DB2 Universal Database
Either DB2 Administration Client or DB2 Universal Database must be installed and configured on the management server.
Next >
Cancel
Help

Figure 7-29 DB choice for the IBM Director

2. Specify a database name as well as an account name used to run the database as shown in Figure 7-30. Either specify an existing user ID or create a new account. In our example we created a new account called itdsdb2. This account has administrator privileges.

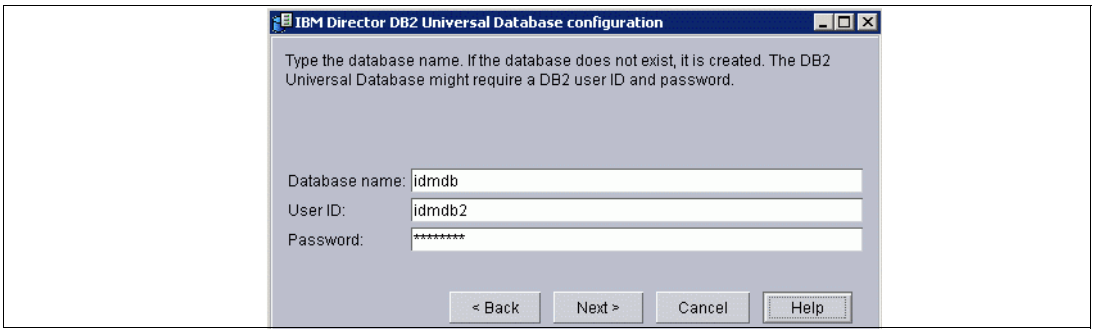
IBM Director DB2 Universal Database configuration
Type the database name. If the database does not exist, it is created. The DB2 Universal Database might require a DB2 user ID and password.
Database name: idmdb
User ID: idmdb2
Password: ********
< Back
Next >
Cancel
Help

Figure 7-30 DB account ID

3. After clicking **Next**, the database is created in the default DB2 instance.

Installing the IBM Director Agent - MSNWIN1.itso.ibm.com

To install the IBM Director Agent, do the following:

1. During the IBM Director installation we also chose IBM Virtualization Engine managed node installation images. This option copies the IBM Director Agent installation images to the C:\Program Files\ibm\VE\ManagedNodes directory on the IBM Director Server.
2. Copy the dir4.20_agent_windows.zip archive to the managed server, expand it, and run the IBMDirectorAgent4.20.exe setup file.
3. We used the default settings. When the installation is complete, restart your server. Two services are created. These services collect local system information and provide a communication interface for the sending this system information to the IBM Director Server. They are the IBM Director Agent WMI CIM Server and the IBM Director Support Program.

7.4 Installing Enterprise Workload Manager

To install the Enterprise Workload Manager (EWLM), we followed these steps:

1. Start the wizard by executing the installvewin.exe file in your Virtualization Engine console staging directory. In our case this is C:\Install_Files\VEC \installvewin.exe. Accept all agreements and choose the appropriate language for the wizard. Then choose the appropriate components. As shown in Figure 7-31, we choose Enterprise Workload Manager and Virtualization Engine managed node installation images.

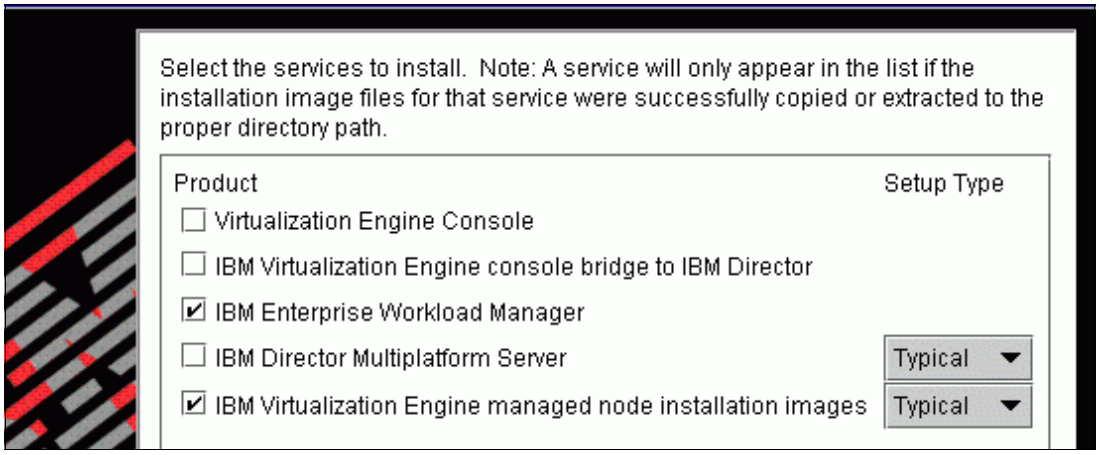
Select the services to install. Note: A service will only appear in the list if the installation image files for that service were successfully copied or extracted to the proper directory path.
Product
Setup Type
Virtualization Engine Console
IBM Virtualization Engine console bridge to IBM Director
IBM Enterprise Workload Manager
IBM Director Multiplatform Server
Typical
IBM Virtualization Engine managed node installation images
Typical

Figure 7-31 Virtualization Engine components choice panel

2. Click **Next** through the remaining screens which summarize and eventually complete installation is complete.

The next sections describe in detail the remaining steps for installing EWLM.

7.4.1 Configuring the EWLM Domain Manager

To configure the EWLM Domain Manager, do the following:

1. Be sure you review the EWLM planning checklist, available at:

   http://publib.boulder.ibm.com/infocenter/eserver/v1r1/en_US/info/ewlminfo/plan.htm

2. After the initial installation completes create a user account for the EWLM WebSphere instance. This account requires administrator privileges.

3. Run the **ConfigWizardDM** command from the C:\Program Files\IBM\VE\EWLM\bin directory. This launches the EWLM configuration GUI as shown in Figure 7-32.

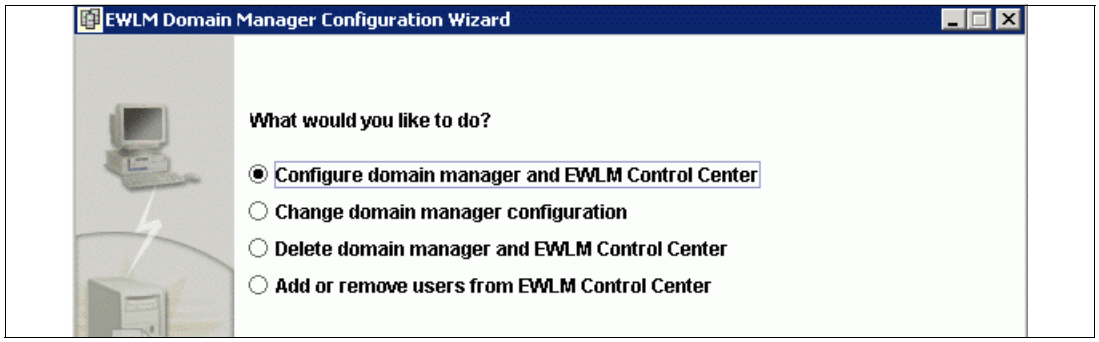
EWLM Domain Manager Configuration Wizard
What would you like to do?
Configure domain manager and EWLM Control Center
Change domain manager configuration
Delete domain manager and EWLM Control Center
Add or remove users from EWLM Control Center

Figure 7-32 The EWLM configuration GUI

4. Choose **Configure domain manager and EWLM Control Center** and click **Next**. Then, specify the working directory for the EWLM Domain Manager. We chose C:\EWLMDM.

   Tip: Place the EWLM working directory in a location that does NOT have spaces (blanks) in the path name.

5. Specify the EWLM Domain Name. We chose itsoewlm. Then specify the ports for running the EWLM Console and WebSphere application server. We manually specified these ports. Make sure the ports are not already in use on your system before specifying them in the domain manager configuration wizard. Figure 7-33 contains the port values that we used in our installation. You can specify any unused ports you wish.

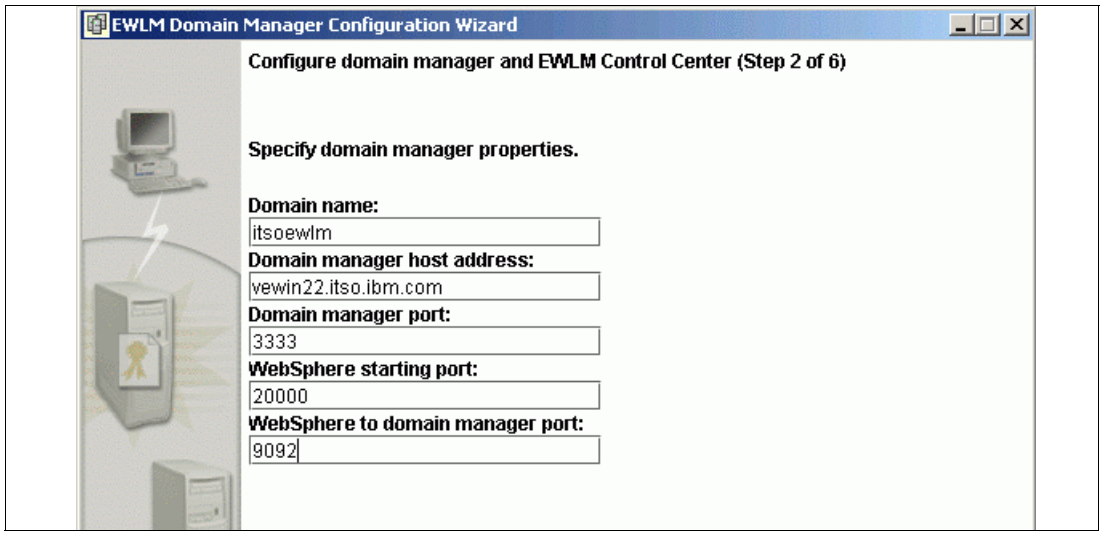

Figure 7-33 Domain manager configuration panel

6. The next screens allow you to configure EWLM security settings. We did not install any of the security features, and instead opted for the default settings.
7. Specify the WebSphere user account that you created earlier, as shown in Figure 7-34.

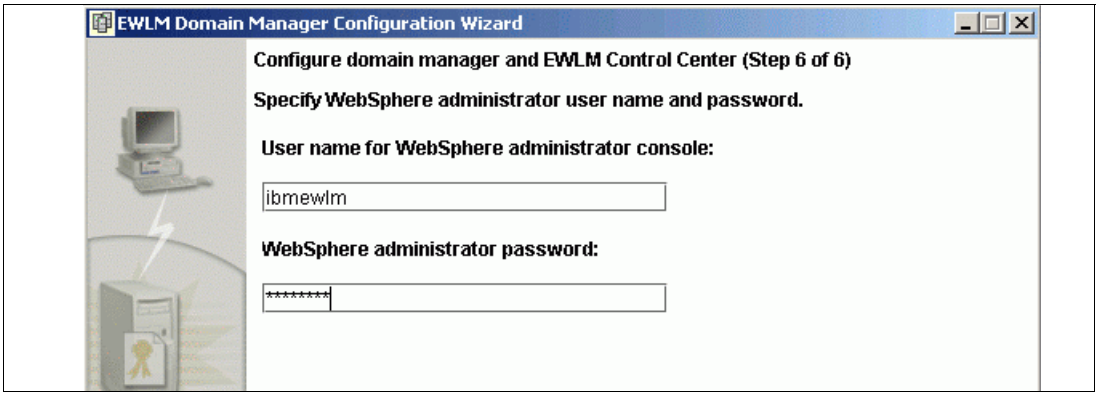
EWLM Domain Manager Configuration Wizard
Configure domain manager and EWLM Control Center (Step 6 of 6)
Specify WebSphere administrator user name and password.
User name for WebSphere administrator console:
ibmewlm
WebSphere administrator password:
****

Figure 7-34 Account parameters

8. Click **Finish** and the configuration begins. A status window describes the configuration progress, and when the configuration it is completed, it displays a summary, as shown in Figure 7-35.

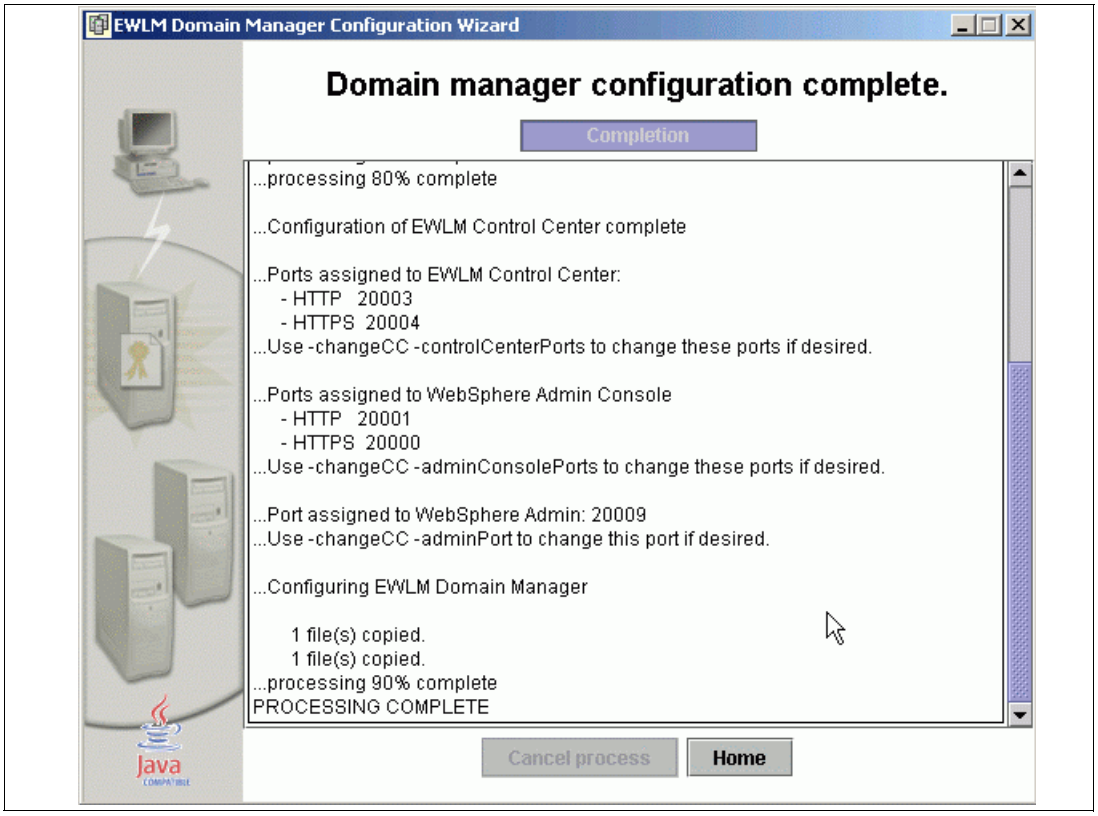
EWLM Domain Manager Configuration Wizard
Domain manager configuration complete.
Completion
...processing 80% complete
...Configuration of EWLM Control Center complete
...Ports assigned to EWLM Control Center:
- HTTP 20003
- HTTPS 20004
...Use -changeCC -controlCenterPorts to change these ports if desired.
...Ports assigned to WebSphere Admin Console
- HTTP 20001
- HTTPS 20000
...Use -changeCC -adminConsolePorts to change these ports if desired.
...Port assigned to WebSphere Admin: 20009
...Use -changeCC -adminPort to change this port if desired.
...Configuring EWLM Domain Manager
1 file(s) copied.
1 file(s) copied.
...processing 90% complete
PROCESSING COMPLETE
Cancel process
Home

Figure 7-35 Final panel of the EWLM Domain Manager configuration

9. Click **Home**. This returns you to the first EWLM configuration window as shown in Figure 7-36. At this point, you create EWLM user accounts.

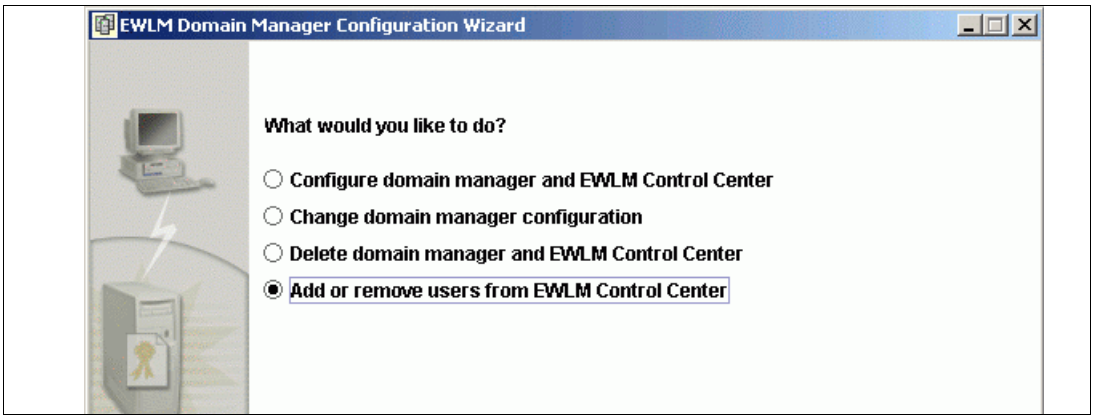
EWLM Domain Manager Configuration Wizard
What would you like to do?
Configure domain manager and EWLM Control Center
Change domain manager configuration
Delete domain manager and EWLM Control Center
Add or remove users from EWLM Control Center

Figure 7-36 EWLM Domain Manager configuration wizard

10. Click **Next**. At the next window, as shown in Figure 7-37, enter the ibmewlm user ID and password that you created earlier. Click **Next** and add users using the Add/Remove users interface. After all accounts are added click **Finish**. You can also create groups using this interface as well. When you are finished, click **Finish**. When the account creation is complete, close the configuration wizard.

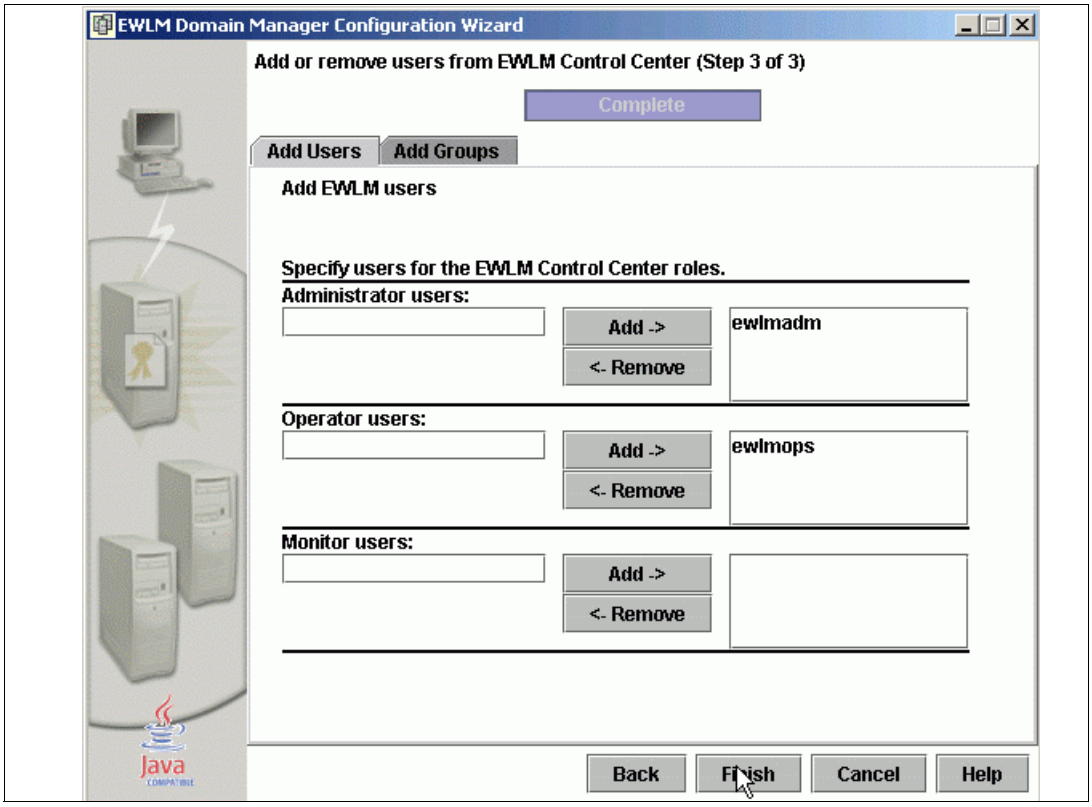

Figure 7-37 Update the account ID parameters

There are command line alternatives to using EWLM Domain Manager Configuration Wizard. These commands are also located in C:\Program Files\IBM\VE\EWLM\bin. The commands are as follows:

- Commands for creating the EWLM domain model:

```
createDM workingDir -adminUser userid -adminPW password -wasPorts port -jp port -ma
address -mp port -dn domainName -auth [None | ServerSSL | ClientServerSSL] -sslks path
-sslpw password
```

- Command for creating EWLM user accounts:

```
changeCC -addUser [workingDir | configID] -adminUser userid -adminPW password -roleUser
user -role [Administrator | Monitor | Operator]
```

7.4.2 Starting EWLM

To start EWLM, follow these steps:

1. Use the **startdm** command to start the EWLM Domain Manager in interactive mode. Alternatively, you can run the domain manager as a Windows service using the **startdmservice** command.
2. Use the **startWAS** command to launch the EWLM Control Center.

Tip: The EWLM WebSphere instance runs in the foreground. Do not log off of the session.

7.4.3 Connecting to EWLM Control Center

To connect to the EWLM Control Center, do the following:

1. Run the **displayCC –ports** command which shows the ports used by the EWLM Control Center as shown in Example 7-4.

Example 7-4 displayCC -ports command display

```
C:\Program Files\IBM\VE\EWLM\bin>displaycc -ports c:\ewlmdm -adminUser ibmewlm
adminPW password
Processing displayCC -ports request.  Please be patient as this may
take a while...
...processing 33% complete
...processing 66% complete
...Ports assigned to EWLM Control Center:
      - HTTP        20003
      - HTTPS       20004
...Use -changeCC -controlCenterPorts to change these ports if desired.
...Ports assigned to WebSphere Admin Console
      - HTTP        20001
      - HTTPS       20000
...Use -changeCC -adminConsolePorts to change these ports if desired.
...Port assigned to WebSphere Admin: 20009
...Use -changeCC -adminPort to change this port if desired.
PROCESSING COMPLETE
C:\Program Files\IBM\VE\EWLM\bin>
```

The output shows that the EWLM Control Center run on ports 20003 (HTTP) and 20004 (HTTPS).

2. Connect to the EWLM Control Center by opening a Web browser to the following:

   http://vewin22.itso.ibm.com:20004/webui/

At this point the EWLM Domain Manager installation is complete.

7.4.4 Installing the EWLM Managed Server

To install the EWLM Managed Server, do the following:

1. During the EWLM installation we also chose IBM Virtualization Engine managed node installation images. This copies the EWLM installation images to the C:\Program Files\ibm\VE\ManagedNodes directory on the EWLM Domain Manager server.
2. Copy the EWLMWinMS.exe file to a temporary directory the managed server. Double click this file to launch the installation.
3. The installation is very straightforward. The default installation directory is C:\Program Files\IBM\VE\EWLMMS.
4. After the installation open a command prompt and run the **`configWizardMS.bat`** command from the C:\Program Files\IBM\VE\EWLMMS\bin directory. This launches the EWLM Managed Server Configuration Wizard GUI.
5. Choose **Configure Managed Server**, and click **Next**.
6. Create a working directory for EWLM Managed Server. We use C:\EWLMMS.

> **Tip:** Place the EWLMMS working directory in a location that does *not* have spaces (blanks) in the path.

7. Specify the name of the EWLM Domain Manager server as well as the domain manager host specified during the EWLM Domain Manager configuration, as shown in Figure 7-38.

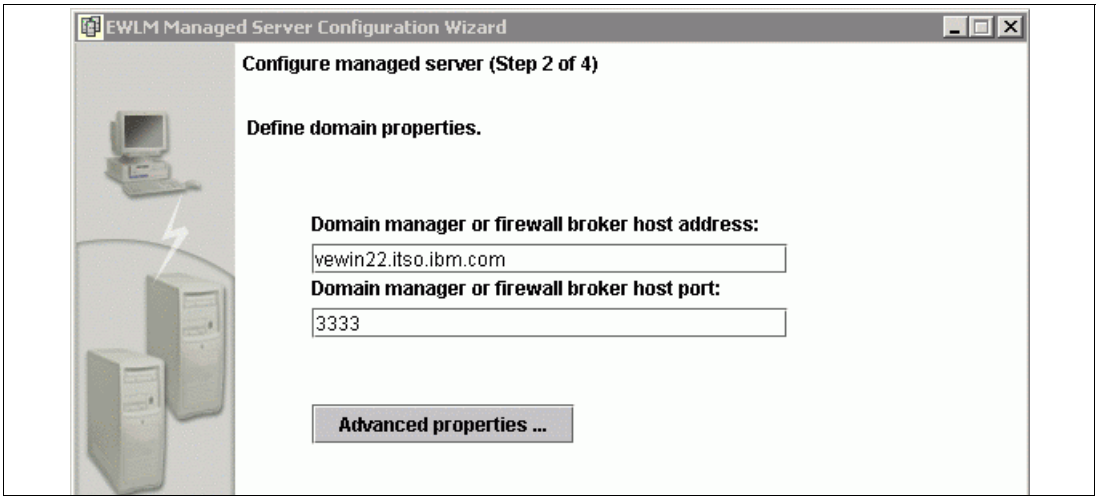
EWLM Managed Server Configuration Wizard
Configure managed server (Step 2 of 4)
Define domain properties.
Domain manager or firewall broker host address:
vewin22.itso.ibm.com
Domain manager or firewall broker host port:
3333
Advanced properties ...

Figure 7-38 Domain manager server configuration

8. Click **Next** twice, and then **Finish**. When the configuration has completed, close the wizard. You can then log to the EWLM Web console.

> **Note:** The EWLM Web console supports Internet Explorer 6.0 with SP1.

The **ewlmWinAdTool** command allows you to verify the applications registered to the ARM interface. Example 7-5 shows the information displayed by such a command.

Example 7-5 ewlmWinAdTool display example

```
Windows Platform Administrative and Diagnostic Tool for EWLM
Ver. 1.00.0130.4230 *
EWLM V1R1.0 (C) Copyright IBM Corporation, 2004.

>>> ARM4PsInfo: 1 process(es) is doing ARM instrumentation...

pid +|- gupid(hex) #ARs #TRs #AIs #TIs #TBs
-------- --- ------------------------ ---- ---- ---- ------- -------
3176 [+] 5D0958A0.01C44DA7.00000C68 : 1/64 2/64 1/256 0/16384 0/16384
AR[00]: E3644F5C00000000h (WebSphere)
TR[00]: E3644FA400000000h (EJB)
.APP: (WebSphere) E3644F5C00000000h
TR[01]: E3644F5C00000000h (WebSphere)
.APP: (IBM DB2 Universal Database) E2D1578100000001h
AI[00]: D794771000000000h (aiid=D794771Oh, aeid=BFE900CB00000003h)
.GRP: (TOT54)
.INS: (TOT54.server1)
.APP: (WebSphere) E3644F5C00000000h
```

Additional information

In addition to the Virtualization Engine Information Center, you can find detailed installation instructions for EWLM in *IBM Enterprise Workload Manager*, SG24-6350.

7.5 Installing Tivoli Provisioning Manager

This section provides information about how to install Tivoli Provisioning Manager.

7.5.1 Installing the prerequisites

You need to install and configure multiple software before installing the Tivoli Provisioning Manager code.

Installing Cygwin

You can download the latest Cygwin distribution from:

http://www.cygwin.com/

At the time of this writing, Cygwin version 1.5.16-1 was available. We chose the list of packages specified in the Tivoli Provisioning Manager 2.1 Installation documentation. Certain packages, OpenSSH for example, select other packages when they are chosen. Leave these add on packages selected. As in the example of OpenSSH, OpenSSL is also selected.

Note: **sh-utils** is part of the Cygwin BASE package group.

Note: We installed **regex** as per *IBM Tivoli Intelligent ThinkDynamic Orchestrator Pre Proof-of-Concept Cookbook for Business Partners* , REDP-3830.

Installing and configuring SNMP

We did not install SNMP on our servers.

Using the prerequisite installer (PICS)

We use the PICS to install Tivoli Directory Server 5.2, DB2 8.1 as well as WebSphere Application Server 5.1. Upon completion each component is patched before we install Tivoli Provisioning Manager. To use the prerequisite installer:

1. Launch the PIC setup executable from the staging directory. You are prompted to first install the PIC on the local system.

Tip: You have to be patient. You are prompted for an installation language, then the location for installing PICS (c:\program files\ibm\pics). The window closes, and then after 10 seconds, another window opens with the actual PICS installation.

2. When it begins, select the products you want PICS to install. We choose the components as shown in Figure 7-39.

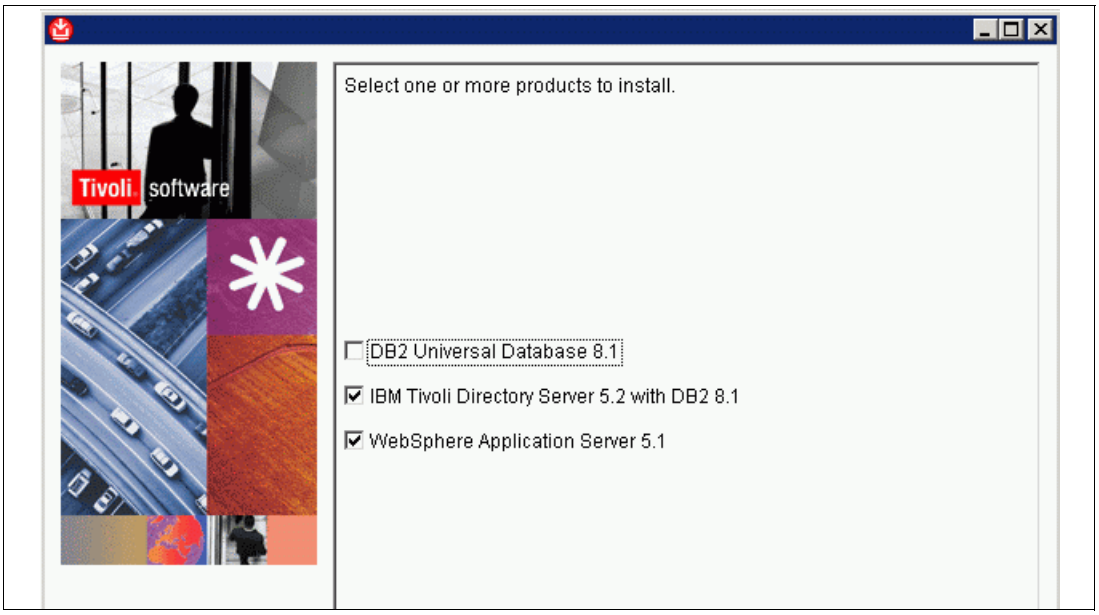
Tivoli software
Select one or more products to install.
DB2 Universal Database 8.1
IBM Tivoli Directory Server 5.2 with DB2 8.1
WebSphere Application Server 5.1

Figure 7-39 Tivoli Provisioning Manager installation wizard selection

3. The first installation step is configuring the DB2 user account. We choose the default.

   Then the installation starts. (As the installation begins, the progress bar does not provide much indication as to what is happening.) DB2 installs first and prompts for the installation file directory (without any change in the progress bar).

   Eventually (after about 10 minutes), another prompt appears requesting the location of the DB2 FP3 directory.

   It takes another 10 minutes for the DB2 Fix Pack 3 installation.

4. You are then prompted for the LDAP installation files. At this point the progress indicator jumps to 33% complete, because technically one third of the products installed by PICS has completed.

   It takes about five minutes for Tivoli Directory Server to install and another prompt appears for WebSphere Application Server 5.1.0. At this point, the progress indicator is as 66%.

5. Follow the prompts for installing WebSphere Application Server. We installed the full version

Installing DB2 license

Start the **DB2licm** command to install the DB2 license in the correct working directory, as shown in Figure 7-40.

```
C:\IBM\SQLLIB\BIN>db2licm C:\Intall_Files\DB28.1\031_WSUE_WIN_32_NLV\db2\license
\db2wsue.lic

DBI1402I License added successfully.


DBI1426I This product is now licensed for use as specified in
         the License Acceptance and License Information
         documents pertaining to the licensed copy of this
         product.  USE OF THE PRODUCT CONSTITUTES ACCEPTANCE OF
         THE TERMS OF THE IBM LICENSE ACCEPTANCE AND LICENSE
         INFORMATION DOCUMENTS, LOCATED IN THE FOLLOWING
         DIRECTORY: "C:\IBM\SQLLIB\license\en"


C:\IBM\SQLLIB\BIN>_
```

Figure 7-40 DB2 licence install

Installing the 10831 Hotfix

We installed the 10831 Hotfix (though not mentioned in the installation documentation). It is included with the generic fixes.

Installing Tivoli Directory Server 5.2

We applied the Tivoli Directory Server 5.2 patch by copying the v3.ibm.at file from the generic fix CD.

Installing the WebSphere Cumulative Fix Pack 3 from generic fix CD

We installed the WebSphere MQ IC38409 patch from the generic fix CD by copying the com.ibm.mqjms.jar and com.ibm.mq.jar files to the appropriate directory. There are multiple so you need to check the MQ environment path in WebSphere Application Server Admin console (wasadmin/wasadmin).

Post prerequisite installation tasks

Note: The Tivoli Provisioning Manager Installation document states that the DB2 and Tivoli Directory Server services need to be started. Tivoli Directory Server is not running and you have to start it. However, DB2 is already running after the PICS completes.

Before Tivoli Directory Server can be started, you must configure a database. If you try to start Tivoli Directory Server, it stops automatically and the following error message displays:

```
The IBM Tivoli Directory Server v5.2 service on Local Computer started and then stopped.
Some services stop automatically if they have no work to do, for example, The
Performance Logs and Alerts service.
```

When the database is configured, this error no longer appears.

Configuring the Tivoli Directory Server and the DB2 database

The DB2 configuration is set up as shown through the Figure 7-41, Figure 7-42, and Figure 7-43 on page 196.

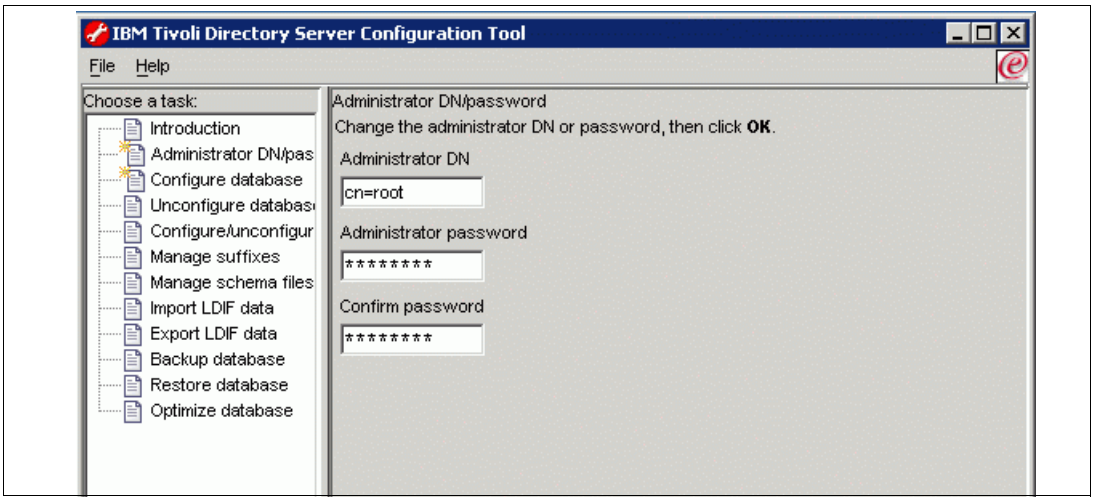

Figure 7-41 Administrator DN account update

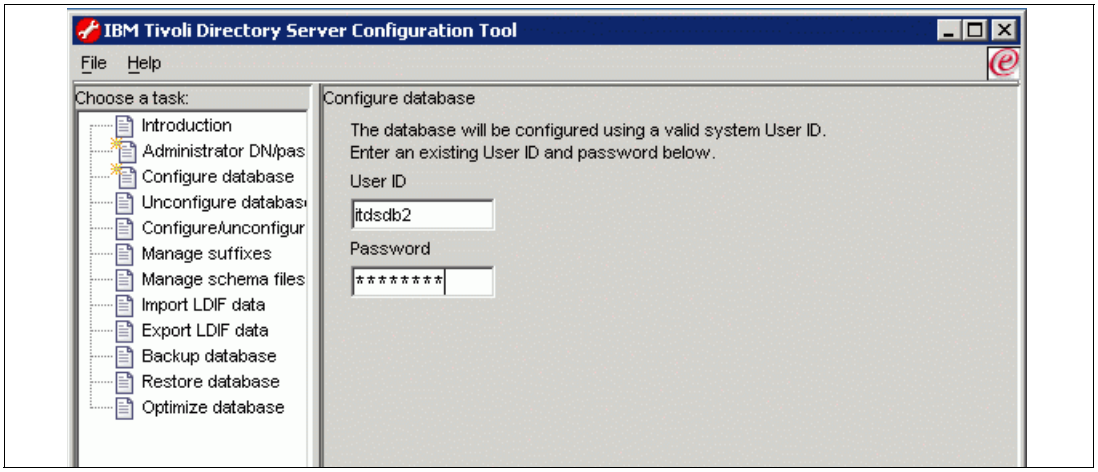

Figure 7-42 DB configuration

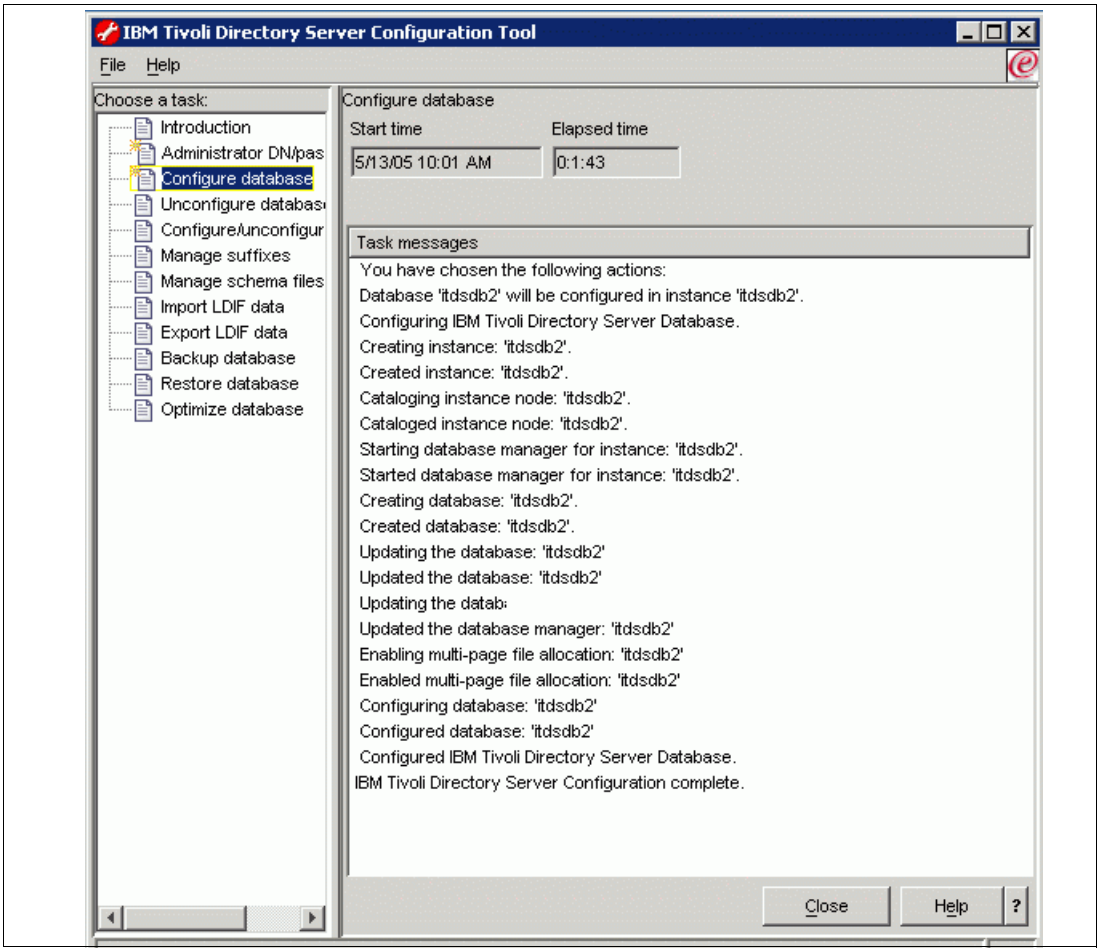

Figure 7-43 DB Configuration summary

Starting Tivoli Directory Server

Tivoli Directory Server takes a while before starting. This also occurs on system startup. A system event ID 7022 is logged stating that Tivoli Directory Server hung on starting. However, the service eventually starts.

Tip: Tivoli Directory Server needs its DB2 database online so it can start. If you configure Tivoli Directory Server for automatic startup, after a reboot you might see an error stating that The IBM Directory Server V5.1 service hung on starting. This occurs because Tivoli Directory Server issues a call to start its database instance and it takes longer than expected for the database to come online. However, when you look at the service in the service control manager, Tivoli Directory Server is in the running state. Ignore the message, or start the service manually. In a production environment, we recommend that you use the dependencies facility to start the services in the order you need.

Important: The Tivoli Provisioning Manager installation creates two services, IBM WebSphere Application Server V5 - IBM_TPM and Tivoli Provisioning Manager.

- The Tivoli Provisioning Manager service is dependant upon the WebSphere service.
- The WebSphere service authenticates via Tivoli Directory Server.

Because Tivoli Directory Server hangs during startup, the WebSphere service will also hang if it is set in automatic startup. This failure takes several minutes before timing out. During this time your system is virtually inaccessible.

You can leave the Tivoli Provisioning Manager and WebSphere Application Server service in manual startup mode, or use the dependencies facility to start the services in the right order.

In a manual mode, after a reboot, log in to the server, confirm that Tivoli Directory Server is started, then manually start the Tivoli Provisioning Manager service. Because of the dependency, starting Tivoli Provisioning Manager causes the WebSphere service to start.

Another alternative is to start Tivoli Provisioning Manager using the interactive tio_start script described in this chapter.

Configuring SSH in Cygwin

To configure SSH in Cygwin, do the following:

1. Create the public key files.
2. Add the public key to authorized_users.
3. Start sshd.
4. Test the public keys by logging into localhost.
5. Copy the public key files to your Tivoli Provisioning Manager managed servers.

Installing Tivoli Provisioning Manager V2.1 code

The installation of the Tivoli Provisioning Manager code follows the documentation that is provided by the product with no specific problems.

Configuring Tivoli Provisioning Manager

After the installation, you need to configure Tivoli Provisioning Manager as described in the following sections.

Building the Tivoli Provisioning Manager data center model

Example 7-6 provides an example of the data center model that we used for our project.

Example 7-6 The Tivoli Provisioning Manager data center model

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- edited with XMLSPY v5 rel. 4 U (http://www.xmlspy.com) by IBM Corporation (IBM
Corporation) -->
<!DOCTYPE datacenter SYSTEM "xmlimport.dtd">
<datacenter>
   <switch-fabric name="ITSO_Fabric" />
   <switch name="ITSO_Switch" locale="en_US" fabric="ITSO_Fabric" failed="false"
in-maintenance="false">
      <switch-module>
         <switch-port failed="false" managed="false" netboot-enabled="false" vlan="1"
port-number="1" enabled="true" layer1-interface-type="unknown" />
         <switch-port failed="false" managed="false" netboot-enabled="false" vlan="1"
port-number="2" enabled="true" layer1-interface-type="unknown" />
```

```
         <switch-port failed="false" managed="false" netboot-enabled="false" vlan="1"
port-number="3" enabled="true" layer1-interface-type="unknown" />
         <switch-port failed="false" managed="false" netboot-enabled="false" vlan="1"
port-number="4" enabled="true" layer1-interface-type="unknown" />
         <switch-port failed="false" managed="false" netboot-enabled="false" vlan="1"
port-number="5" enabled="true" layer1-interface-type="unknown" />
         <switch-port failed="false" managed="false" netboot-enabled="false" vlan="1"
port-number="6" enabled="true" layer1-interface-type="unknown" />
         <switch-port failed="false" managed="false" netboot-enabled="false" vlan="1"
port-number="7" enabled="true" layer1-interface-type="unknown" />
         <switch-port failed="false" managed="false" netboot-enabled="false" vlan="1"
port-number="8" enabled="true" layer1-interface-type="unknown" />
      </switch-module>
      <switch-vlan vlan-number="1" />
   </switch>
   <subnetwork name="ITSOVE" ipaddress="10.1.1.0" netmask="255.255.255.0">
      <vlan vlan-number="1" fabric="ITSO_Fabric" />
   </subnetwork>
   <software-category name="Windows 2003" />
   <software name="Windows 2003 Server" locale="en_US" type="Operating System"
version="5.2.3790 SP1" package_path="package" category="Windows 2003" />
   <software-category name="TIO" />
   <software name="TIO" locale="en_US" type="Software" version="2.1" package_path="package"
category="TIO" />

   <spare-pool name="ITSOVE Windows" locale="en_US" os-type=" Windows" vlan="1"
fabric="ITSO_Fabric">
      <server name="vewin23" locale="en_US">
         <nic name="Local Area Connection" macaddress="000D609C3AE4"
netboot-enabled="false" managed="true" connected-to-switch="ITSO_Switch"
connected-to-port="1">
            <network-interface name="Local Area Connection" ipaddress="10.1.1.180"
netmask="255.255.255.0" />
         </nic>
         <property name="platform" value="IBM eServer HS20"/>
         <sap name="scp client" port="22" host="false" app-protocol="SCP"
protocol-type="ipv4" locale="en_US">
            <credentials search-key="primary" is-default="true">
               <password-credentials username="root"/>
            </credentials>
         </sap>
         <sap name="ssh service" port="22" host="true" app-protocol="SSH"
is-device-model="SSH Service Access Point" protocol-type="ipv4" locale="en_US">
            <default-sap operation-type="execute-command"/>
            <credentials search-key="secondary" is-default="false">
               <password-credentials username="tioadmin" password="password"/>
            </credentials>
            <credentials search-key="primary" is-default="true">
               <rsa-credentials username="tioadmin"/>
            </credentials>
         </sap>
         <sap name="snmp-set" port="161" host="true" app-protocol="SNMP"
is-device-model="SNMP V1 Service Access Point" protocol-type="ipv4" context="write"
locale="en_US">
            <default-sap operation-type="set-attribute"/>
            <credentials search-key="primary" is-default="true">
               <snmp-credentials community="public"/>
            </credentials>
         </sap>
```

```
         <sap name="snmp-get" port="161" host="true" app-protocol="SNMP"
is-device-model="SNMP V1 Service Access Point" protocol-type="ipv4" context="read"
locale="en_US">
            <default-sap operation-type="get-attribute"/>
            <credentials search-key="primary" is-default="true">
               <snmp-credentials community="public"/>
            </credentials>
         </sap>
         <software-association software="Windows 2003 Server" state="running"/>
         <software-association software="TIO" state="running"/>
      </server>
   </spare-pool>
   <kanaha-config>
      <dcm-object id="0">
         <property component="DEPLOYMENT_ENGINE" name="PackageIntermediateHost"
value="vewin23"/>
         <property component="DEPLOYMENT_ENGINE" name="PackageRepositoryHost"
value="vewin23"/>
         <property component="DEPLOYMENT_ENGINE" name="PackageWebHost" value="vewin23"/>
         <property component="DEPLOYMENT_ENGINE" name="SoftwareRepositoryHost"
value="vewin23"/>
      </dcm-object>
   </kanaha-config>
</datacenter>
```

Re-initializing the Tivoli Provisioning Manager database

The default data center model is a sample file called venice.xml. This is the default data center model that is created during the Tivoli Provisioning Manager installation. This information is useful for seeing how a fully populated data center model looks. However, before configuring and using Tivoli Provisioning Manager we recommend removing the venice.xml data and start with a fresh data center model.

In order to remove this data center model data populated by the venice.xml file you need to re-initialize the Tivoli Provisioning Manager database.

Tip: We recommend starting fresh, even for testing purposes. The default Tivoli Provisioning Manager data center values contain many dependencies which can be hard to remove if you are not familiar with the interface**.**

You can re-initialize the Tivoli Provisioning Manager database by doing the following:

- Importing the data center model by running the **`xmlimport`** command and specifying the data center model that you created.
- Manually adding data center devices using the GUI, which provides an alternative to importing an XML file.

For more information, see the following Technote:

http://www-1.ibm.com/support/docview.wss?uid=swg21190846

Testing the Tivoli Provisioning Manager installation

To test the installation, you can run a workflow to test SSH connectivity using the SSH **`execute`** command.

8

Installing and configuring the Virtualization Engine for IT management in an OS/400 environment

This chapter demonstrates the Virtualization Engine solution being integrated into an iSeries environment. The primary systems services are running on IBM iSeries hardware. This scenario is for example purposes only and does not imply a recommendation for your environment.

We did not implement this scenario in the our project because of time dependency. However we want to demonstrate that the IBM @server iSeries is definitely part of the Virtualization Engine initiative. You can find this scenario at:

http://publib.boulder.ibm.com/infocenter/eserver/v1r1/en_US/index.htm?info/icmain.htm

© Copyright IBM Corp. 2005. All rights reserved.

8.1 Scenario description

In this scenario, MyITSOCo, Inc. wants to create an enterprise-wide solution for managing their IT environment, which consists primarily of OS/400 systems and a few Windows systems. Administrators in the current environment use the monitoring tools available in iSeries Navigator Management Central and other products to help manage their systems as easily and centrally as possible. MyITSOCo, Inc. is currently using OS/400 on iSeries hardware to host a custom application that they use for online transaction processing.

The company is interested in the Virtualization Engine as a means of centralizing systems management to make it easier for them to monitor and optimize system resources.

8.1.1 Objectives

As the lead administrator in MyITSOCo, Inc., your goal is to fully utilize the capabilities of the Virtualization Engine, including the use of systems provisioning and monitoring. To take advantage of these capabilities, you are installing and configuring Virtualization Engine systems services within your IT environment.

The goals of this scenario is:

- To install Enterprise Workload Manager (EWLM) to monitor application transactions across operating systems.
- To plan the installation of Tivoli Provisioning Manager to provision resources.
- To install IBM Director Multiplatform to monitor system status and manage enabled IBM hardware.
- To install the Virtualization Engine console to use IBM Director, Tivoli Provisioning Manager, and EWLM interfaces from one central location.

8.1.2 Topology

Figure 8-1 shows the company's network topology and where the systems administrator decided to implement the Virtualization Engine systems services.

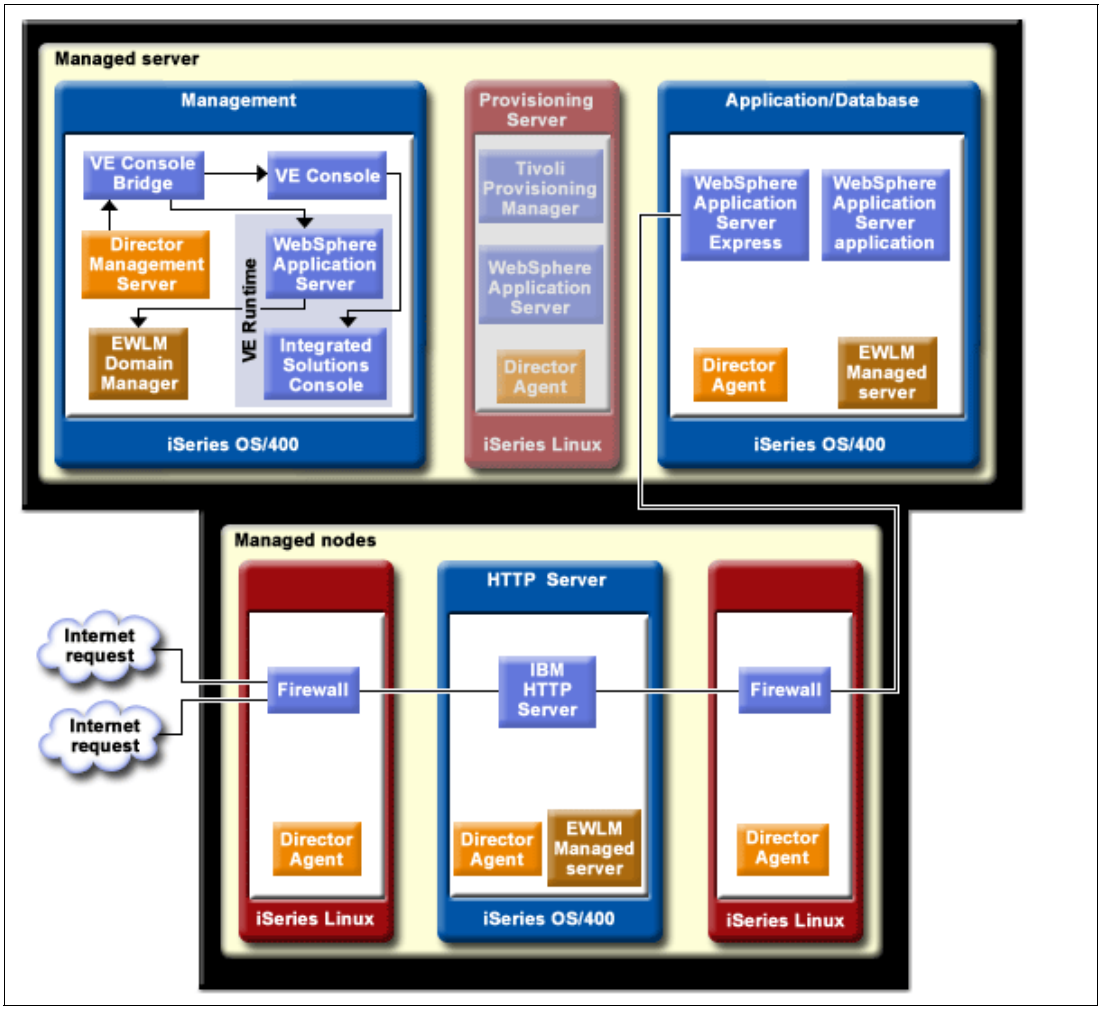

Figure 8-1 iSeries scenario topology

Prerequisites and assumptions

Successful implementation of this scenario requires that the following prerequisites are met:

- Print and record the Virtualization Engine planning advisor results before beginning the installation process.
- Ensure that all systems for which you are installing the Virtualization Engine upon meet the minimum system requirements.
- The administrator who performs the installation must have *ALLOBJ, *SECADM, *JOBCTL, and *IOSYSCFG special authorities.
- The administrator has read the LDAP primer information to ensure a basic understanding of the LDAP principles needed to configure a Directory Server.
- This scenario assumes that MyITSOCo, Inc. is currently is using a Directory Server on OS/400 and that they are using it with the predefined basic configuration for the directory server. They do not use this directory server for any other applications. Therefore, this scenario does not describe how MyITSOCo configures the IBM Directory Server 5.1.

Rather, it describes what values are needed as part of their overall Virtualization Engine installation and configuration process.

- This scenario does not have MyITSOCo, Inc. setting up systems provisioning at this time. At this point in their implementation plans, MyITSOCo, Inc. will set up the main Virtualization Engine environment and systems services and begin planning their future implementation of systems provisioning.

8.2 Configuration steps

To configure the complete Virtualization Engine environment that this scenario describes, complete the following tasks:

- Task 1: Completing and printing the planning worksheets
- Task 2: Completing the LDAP installation for the Virtualization Engine console
- Task 3: Installing the Virtualization Engine
- Task 4: Installing and configuring the Virtualization Engine agents on managed servers
- Task 5: Configuring EWLM Domain Manager
- Task 6: Installing and configuring EWLM Managed Servers
- Task 7: Beginning the planning process for implementing systems provisioning

8.2.1 Task 1: Completing and printing the planning worksheets

To ensure a successful Virtualization Engine implementation, complete the Virtualization Engine planning process.

MyITSOCo, Inc. is installing EWLM, the Virtualization Engine console, and IBM Director server for Virtualization Engine management services. They are also installing a number of IBM Director Agents and EWLM Managed Servers on systems throughout their enterprise so that they can use the capabilities of the Virtualization Engine systems services to manage them.

To install the Virtualization Engine console and EWLM on their OS/400 Management Server, MyITSOCo, Inc. will use the Virtualization Engine Installation Wizard. They will install the IBM Director and the managed servers and agents throughout their environment by using the appropriate, separate installation processes.

The administrative team worked through the Virtualization Engine planning advisor and other planning documentation to plan their overall installation. The administrative team has also reviewed the Virtualization Engine console LDAP primer to ensure they understand the LDAP principles needed for the installation. Before MyITSOCo, Inc. begins the physical installation process, they complete and print the following worksheets as a guide to the installation:

- LDAP Server Connection Information worksheet, which is located in the Installing the Virtualization Engine console topic.
- EWLM Domain Manager planning worksheet.
- EWLM Managed Server planning worksheet .
- EWLM firewall broker planning worksheet (only necessary if you are installing a firewall broker).
- Planning checklist for IBM Director on OS/400.
- Planning checklist for IBM Director Agent on Linux.
- Virtualization Engine console bridge for IBM Director installation walkthrough checklist.

8.2.2 Task 2: Completing the LDAP installation for the Virtualization Engine console

This section includes the steps that MyITSOCo, Inc. needs to perform to complete the LDAP installation.

Step 1: Installing IBM Directory Server 5.1

The IBM Directory Server is installed automatically when you install OS/400. This installation includes a default configuration. The directory server uses the default configuration when all of the following are true:

- Administrators have not run the IBM Directory Server Configuration Wizard or changed directory settings with the properties pages.
- IBM Directory Server publishing is not configured.
- The IBM Directory Server cannot find any LDAP DNS information.

If the IBM Directory Server uses the default configuration, then the following occur:

- The IBM Directory Server starts automatically when TCP/IP starts.
- The system creates a default administrator, cn=Administrator. It also generates a password that is used internally. If you need to use an administrator password later, you can set a new one from the IBM Directory Server property page.
- A default suffix is created that is based on the system's IP name. A system objects' suffix is also created based on the system name.
- The IBM Directory Server uses the default data library QUSRDIRDB. The system creates it in the system ASP.
- The server uses port 389 for non-secure communications. If a digital certificate has been configured for LDAP, secure sockets layer (SSL) is enabled and port 636 is used for secure communications.

Step 2: Configuring the IBM Directory Server 5.1 server

Note: You must have *ALLOBJ and *IOSYSCFG special authorities to configure the server.

To configure the IBM Directory Server 5.1 server, follow these steps:

1. In iSeries Navigator click **Network** → **Servers** → **TCP/IP**.
2. Click **Configure system as Directory server** in the Server Configuration tasks window at the bottom right of iSeries Navigator.

   The Directory Server Configuration Wizard appears.
3. Click **Configure a local LDAP directory server** on the IBM Directory Server Configuration Wizard - Welcome window.
4. Click **Next** on the IBM Directory Server Configuration Wizard - Welcome window.
5. Select **No** on the IBM Directory Server Configuration Wizard - Specify Settings window. This allows you to configure the LDAP server without the default settings. Click **Next**.
6. Select **Next** to keep the default library disk pool on the Select Directory Storage Location window.

7. Deselect **System-generated** on the IBM Directory Server Configuration Wizard - Specify Administrator DN window and enter the following:

Administrator DN	cn=name
Password	your password here
Confirm password	your password here

8. Click **Next** on the IBM Directory Server Configuration Wizard - Specify Administrator DN window.
9. Enter the suffix that you want to add in the Suffix DN field and click **Add**. In this example, the administrator added the suffix o=ibm, c=us.
10. Select **Yes**, use all IP addresses on the IBM Directory Server Configuration Wizard - Select IP Addresses window. Click **Next**.
11. Select **Yes** on the IBM Directory Server Configuration Wizard - Specify TCP/IP Preference window. Click **Next**.
12. On the IBM Directory Server Configuration Wizard - Specify Information to Publish window, leave the defaults and click **Next**.
13. On the IBM Directory Server Configuration Wizard - Specify Publishing DN window, select the Suffix that you added in previously and then click **Next**. In this example, this would be o=ibm, c=us.
14. Click **Finish** on the IBM Directory Server Configuration Wizard - Summary window.
15. In iSeries Navigator, click **Network** → **Servers** → **TCP/IP**. Right-click **Directory** and select the start option to start the Directory Server.

Note: You might need to refresh the screen to see the started status.

Step 3: Setting up the Web Configuration Tool

To do the initial setup of the Web Configuration Tool, do the following:

1. Install IBM WebSphere Application Server - Express (5722-IWE Base and Option 2) and associated prerequisite software if they are not already installed.
2. Enable the system application server instance in the HTTP ADMIN server instance by doing the following:
   a. Start the HTTP ADMIN server instance by doing one of the following.
      i. In iSeries Navigator, click **Network** → **Servers** → **TCP/IP** and right-click **HTTP Administration**. Then click **Start**.
      ii. On an OS/400 command line, enter `STRTCPSVR SERVER(*HTTP) HTTPSVR(*ADMIN)`.
   b. Log on to the IBM Web Administration for iSeries. Use an OS/400 user profile and password to log on to the iSeries Tasks page (http://your_server:2001), then click **IBM Web Administration for iSeries**.
   c. From the HTTP Server Administration your_server page, go to the Manage tab and then the HTTP Servers tab. Make sure ADMIN - Apache is selected in the Server drop-down list. From the options in the left pane of the page, click **General Server Configuration**.

      Note: You might need to expand the section Server Properties in order to see the General Server Configuration option.

   d. Select **Yes** to start the system application server instance when the Admin server is started. Click **OK**.

3. Set the WebSphere Application Server to use the SYSINST by doing the following:
   a. Click **WebSphere Application Server** from the left pane options.
   b. Select **WebSphere Application Server - Express 5.0**.
   c. From the WebSphere instance drop-down list, select **SYSINST**.

      Note: If SYSINST is not present in the drop-down list, restart the ADMIN server.

   d. From the Start all WebSphere application server(s)... drop-down list, select **Yes**.
   e. From the Stop all WebSphere application server(s)... drop-down list, select **Yes**. Click **OK**.
4. Restart the HTTP ADMIN server instance by clicking **Restart** (under the HTTP Servers tab).

   You can also stop and start the HTTP ADMIN server instance using the iSeries Navigator or an OS/400 command line. You can stop the HTTP ADMIN server instance by doing one of the following:
   a. In iSeries Navigator click **Network** → **Servers** → **TCP/IP** and right-click **HTTP Administration**. Then click **Stop**.
   b. On an OS/400 command line, enter `ENDTCPSVR SERVER(*HTTP) HTTPSVR(*ADMIN)`.

   You can start the HTTP ADMIN server instance by doing one of the following:
   a. In iSeries Navigator click **Network** → **Servers** → **TCP/IP** and right-click **HTTP Administration**. Then click **Start**.
   b. On an OS/400 command line, enter `STRTCPSVR SERVER(*HTTP) HTTPSVR(*ADMIN)`.
5. Log on to the Directory Server Web Administration Tool.
   a. Bring up the Login page by doing one of the following:
      i. From iSeries Navigator, select your server and click **Network** → **Servers** → **TCP/IP**, right-click IBM Directory Server, and click Server Administration.
      ii. From the iSeries Tasks page (http://your_server:2001), click **IBM Directory Server for iSeries**.
      iii. From a Web browser and enter the following address:

      `http://localhost:2001/IDSWebApp/IDSjsp/Login.jsp`
   b. Select **Console Admin** in the LDAP Hostname field.
   c. Enter `superadmin` in the Username field.
   d. Enter `secret` in the Password field.
   e. Click **Login**. The IBM Directory Server Web Administration Tool page is displayed.
6. Change the console administration login (optional) by doing the following:
   a. Click **Console administration** in the left pane to expand the section, and then click **Change console administrator login**.
   b. Type a new console administration login name in the Console administrator login field.
   c. Type the current password (secret) in the Current password field.
   d. Click **OK**.

7. Change the console administration password (optional) by clicking **Change console administrator password** in the left pane.

Step 4: Configuring the Virtualization Engine console LDAP settings using the IBM Directory Server Web Administration Tool

After you have completed the previous steps and you have the Web Administration Tool up and running, you can set up the directory server to store the Virtualization Engine console LDAP settings that are required for the Virtualization Engine console installation. To do this:

1. Start the Directory Server by following these steps:
   a. In iSeries Navigator, select **Network** → **Servers** → **TCP/IP**.
   b. Right-click Directory and select Start.

   Note: The Directory Server can take several minutes to start, depending on the speed of your server and the amount of available memory. The first time you start the Directory Server might take several minutes longer than usual because the server must create new files. Similarly, when starting the Directory Server for the first time after upgrading from an earlier version, it can take several minutes longer than usual because the server must migrate files.

   You can also start the Directory Server from the character-base interface using the **STRTCPSVR *DIRSRV** command. Additionally, if the Directory Server is configured to start when TCP/IP starts, you can also start it using the **STRTCP** command.

   You can start the Directory Server in configuration only mode from the character-base interface by using the **TRCTCPAPP APP(*DIRSRV) ARGLIST(SAFEMODE)** command. Configuration only mode starts the server with only the cn=configuration suffix active and does not depend on successful initialization of the database backends.

2. Log in to the console by doing the following:
   a. Open a Web browser and type the following address:

      `http://localhost:2001/IDSWebApp/IDSjsp/Login.jsp`

      The IBM Directory Server Web Administration login page panel is displayed.
   b. Log in as Console Admin, the default selection in the LDAP Hostname field.
   c. In the Username field, enter `superadmin`.
   d. In the Password field, enter `secret`.
   e. Click **Login**
3. Add your LDAP server to the console by doing the following:
   a. Expand Console administration in the navigation area.
   b. Click **Manage console servers**.
   c. Click **Add**.
   d. Enter the host name address. For example `myserver.mycity.mycompany.com`.
   e. Specify the port numbers or accept the defaults.
   f. The Virtualization Engine console does not support SSL connections between the LDAP server. Do not enable SSL.

g. Click **OK** to apply the changes or click **Cancel** to exit the panel without making any changes.

h. When you have finished setting up the console, click **Logout** to exit.

4. Log in to your LDAP server by doing the following:

   a. At the Logout successful panel that is displayed when you exit Console administrator, you see the following message: `If you have been accidentally logged out then you need to log in again by clicking here`. Click the word *here* to go to the IBM Directory Server Web Administration Login Page.

   b. At the IBM Directory Server Web Administration login page, select the LDAP host name for your machine.

   c. Enter the domain name and the password for that server. This is the domain name that you set up during the server installation process. The format should be `cn=DN user ID`.

   d. Click **Login**.

   You are now logged in to your LDAP machine and ready to configure the necessary Virtualization Engine console settings through the Web Administration Tool.

5. Create the suffix DN entry by following these steps:

   Note: If the domain name entry exists, you can skip this step.

   a. Expand Directory Management in the navigation area.

   b. Click **Add an entry**.

   c. Select a structural object class that is appropriate for the suffix. For example, if the suffix DN that you added during the server installation process was `o=ibm,c=us`, then the appropriate object class would be organization. Use the following list to determine which object class is appropriate for your suffix entry:

      - Suffix DN Object class
      - dc= domain
      - o= organization
      - ou= organizationalUnit
      - c= country

   d. Click **Next**.

   e. An auxiliary object class is not required. Click **Next** from the auxiliary object classes panel to skip to the next panel.

   f. Split the suffix DN into a relative DN (first component of the suffix DN) and parent DN (remainder of the name). For example, if your suffix DN is `o=ibm,c=us` then the relative DN is `o=ibm` and the parent DN is `c=us`.

   g. Enter any required attributes. In this case the attribute used in the name is typically the only required attribute, and the Required Attributes can be left empty.

   h. Click **Finish**.

6. Create cn=users and cn=groups container objects as follows:

   a. Expand Directory Management in the navigation area.

   b. Click **Manage entries**.

   c. Select the suffix DN that you created in previously and click **Add**.

   d. Select the container for the Structural object class and click **Next**.

e. An auxiliary object class is not required. Click **Next** from the auxiliary object classes panel to skip to the next panel.

f. Enter `cn=users` (or `cn=groups`) in the Relative DN field.

g. The Parent DN field should contain the value of the suffix DN to which you are adding the cn=users container. For example, if your Suffix DN is `o=ibm,c=us` then this is the value that should appear in the Parent DN field.

h. You do not need to enter any required attributes. In this case the cn attribute used in the name is the only required attribute, and the Required Attributes can be left empty.

i. Click **Finish**.

j. To create a cn=groups container object, repeat these steps and substitute `cn=groups` in place of `cn=users`.

7. Add users by following these steps:

a. Expand Directory Management in the navigation area.

b. Click **Manage entries**.

c. Select the suffix DN that you created previously and click **Expand**.

d. Select the cn=users entry that you created previously and click **Add**.

e. Select **inetOrgPerson** for the Structural object class and click **Next**.

f. An auxiliary object class is not required. Click **Next** from the auxiliary object classes panel to skip to the next panel.

g. Enter the user ID that you will use to sign into the Virtualization Engine console in the Relative DN field (for example `uid=iscadmin`).

h. The parent DN field should contain the value of the cn=users container and the suffix DN to which you are adding the user entry. For example, if your suffix DN is `o=ibm,c=us`, then `cn=users,o=ibm,c=us` is the value that should appear in the parent DN field.

i. Under Required Attributes, enter the common name that is associated with the relative DN in the cn field. For example, if your relative DN is `uid=iscadmin`, then you would enter `iscadmin` in the cn field.

j. Under Required Attributes, enter the value of the relative DN in the sn field. For example, if your relative DN is `uid=iscadmin`, then you would enter `iscadmin` in the sn field.

k. Go to the Other attributes tab.

l. Scroll down to the userPassword field and enter the password that is associated with the user ID.

m. Click **Finish**

n. Verify that you can authenticate as that user by starting another IBM Directory Server Web Administration Tool session and logging in as the user that you just created You must log in with the fully qualified userID (for example: `uid=iscadmin,cn=users,o=ibm,c=us`).

8. Add the Portal administrator group by following these steps:

a. Expand Directory Management in the navigation area.

b. Click **Manage entries**.

c. Select the suffix DN that you created previously and click **Expand**.

d. Select the cn=groups entry that you created previously and click **Add**.

e. Select **groupOfUniqueNames** or **groupOfNames** for the Structural object class and click **Next**.

f. An auxiliary object class is not required. Click **Next** from the auxiliary object classes panel to skip to the next panel.

g. Enter the Portal administrator group in the relative DN field (for example `cn=iscadmins`).

h. The parent DN field should contain the value of the cn=groups container and the suffix DN to which you are adding the user entry. For example, if your suffix DN is `o=ibm,c=us`, then `cn=groups,o=ibm,c=us` is the value that should appear in the parent DN field.

i. Under Required Attributes, enter the common name associated with the relative DN in the cn field. For example, if your relative DN is `cn=iscadmins`, then you would enter `iscadmins` in the cn field.

j. Under Required Attributes, enter the DN of the user ID that is associated with the group in the uniqueMember field (for example, `uid=iscadmin,cn=users,o=ibm, c=us`).

k. Click **Finish**.

8.2.3 Task 3: Installing the Virtualization Engine

The first phase for installing the Virtualization Engine takes place on the system that you designated as your Virtualization Engine management server during the planning process. For MyITSOCo, this is Management Server 1, an iSeries OS/400 system with a host name of mgmtsvr1. When you use the physical media for installation, the installation process consists of two wizards: a Media Transfer wizard and an Installation wizard. The two wizards appear in sequence automatically and guide you through the installation process.

Note: If you have worked through the Virtualization Engine planning advisor and other planning documentation to plan your overall installation, you might prefer to download the Virtualization Engine media that you need from the Passport Advantage IBM Web site:

`http://www.ibm.com/software/passportadvantage`

You can then decompress the files on a network file system so that you can use the files on any system that can access this file system. Alternatively, you can use FTP in binary mode to transfer the files to each installation system that needs them. When you have the files available to the installation system, you can launch the installVEOS400.exe file from a Windows system on your network to start the Virtualization Engine Installation wizard. Then, start with “Step 2: Running the Virtualization Engine Installation wizard on Management Server 1” on page 212.

Step 1: Running the Virtualization Engine Media Transfer wizard on Management Server 1

To run the wizard, follow these steps:

1. Insert the CD with the label START HERE - Installation CD (1 of 4) Virtualization Engine console and Runtime for OS/400.

Note: If the installation wizard does not appear, you might need to launch the CD from the appropriate CD device.

2. When the wizard displays, review the Welcome panel for the media transfer process and click **Next**.
3. On the Copy Location for Media Transfer window, specify where to copy the Virtualization Engine CDs, or click Browse to select a location. The default location is C:\ProgramFiles\IBM\VE\Install\OS400. The MyITSOCo administrator accepts the default location. After the wizard checks system requirements, the contents of the CD are copied to the location that you specified.
4. On the Launch Virtualization Engine Installation Assistant, select **Yes**.
5. Click **Next** to start this browser-based tool to help you determine which CDs you need to copy for installing Virtualization Engine systems services for this server. Print the list of media pack installation CDs that the installation assistant displays and close the installation assistant. Alternatively, you can leave the installation assistant open so that you can refer back to the list of CDs and have the Documentation page available to you during installation.
6. Return to the Begin Media Transfer window, select **Begin copying CDs**, and click **Next**.
7. Insert the first CD to copy into the drive and click **Next** to begin copying the CD. When the copy process completes, the Continue CD Copy page displays.
8. Select **Yes** and click **Next**. Use your printed list as a guide for which CDs to copy and repeat this step until you have copied all the needed CDs.
9. After you copy all CDs, select **No, all the CDs have been copied** and click **Next**.
10. On the Begin Virtualization Engine Installation page, click **Next** to start the Virtualization Engine Installation wizard.
11. Review and accept the license agreement in your preferred language and click **Next**.

Step 2: Running the Virtualization Engine Installation wizard on Management Server 1

After you have copied all the required media or unzipped the downloaded files, you can use the Virtualization Engine Installation wizard to guide you through the process of installing your Virtualization Engine systems services. As you begin the wizard, have your completed planning worksheets and checklists available so that you can complete the wizard pages quickly and correctly.

Complete these steps to use the installation wizard:

1. When the Welcome page displays, click **Next**.
2. Read the Software License Agreement, select I accept the terms of the license agreement, and click **Next**.
3. On the Virtualization Engine Install Location page, note that the installation path is /QIBM/ProdData/VE.

   Note: If you downloaded the files for installation, you might need to create the /QIBM/ProdData/VE directory.

4. On the Select Services to Install page, select the following systems services and images to install:
   - Virtualization Engine console
   - Virtualization Engine console bridge for IBM Director
   - EWLM

– Virtualization Engine managed server installation images as follows:

i. Select **Typical** for EWLM. This changes the setting to Custom, which allows you to select only those EWLM features that you need for the installation.

ii. Select **Typical** for Virtualization Engine managed server installation images. This changes the setting to Custom and allows you to select only those EWLM Managed Server images that you need for the installation.

iii. Click **Next**.

iv. On the Select Features for EWLM page, deselect **Single System** and click **Next**. Only the EWLM Domain Manager is selected for installation.

v. On the Select Features for managed nodes page, deselect all images and agents that you do not intend to use for installation on your systems and click **Next**. For this scenario, we leave the following selected: IBM Director Agent for OS/400, IBM Director Agent for Linux (64 bit), or IBM EWLM Managed Server for OS/400.

5. Review the two Software License Agreements that display, select I accept the terms of the license agreement, and click **Next**.

Note: At this point, the wizard requests information that it needs to install the ISC. The Virtualization Engine console runs within the ISC. The wizard also requests the information that it needs to install a WebSphere Application Server for use by the ISC.

6. Read the information and click **Next**.
7. In the Configuration of WebSphere Application Server Installation page, specify the following to configure basic server settings:
   – In the Node name field, enter descriptive text to identity the WebSphere Application Server instance. Alternatively, you can verify and accept the default name. For this scenario, we enter `WASforVEC`.
   – In the Hostname or IP address field, specify either the fully-qualified host name or the IP address of this server. Alternatively, you can verify and accept the default host name. For this scenario, we enter `mgmtsvr1.myitsoco.com`.
8. Review the information about the LDAP Settings page, and click **Next**.

Note: Use your completed LDAP Server Connection Information worksheet as a reference to help you enter the required LDAP values in the next steps.

9. In the LDAP Server Connection Information page, specify the following information:
   a. In the LDAP server host field, specify the fully-qualified host name for your IBM Directory Server. This is the Directory Server that the Virtualization Engine console and ISC use. Alternatively, you can verify and accept the default host name. For our scenario, we enter `mgmtsrv1.myitsoco.com`, which is the host name of the current system and is the existing directory server location.
   b. In the LDAP server port field, specify the port number for the LDAP server. Alternatively, verify and accept the default port number 389, as we did for this scenario.
   c. In the DN used for directory updates field, specify the LDAP administrator user identity, in the form of an LDAP distinguished name. This is the user ID of the LDAP directory administrator. Alternatively, verify and accept the default DN. In this scenario, we enter `cn=Administrator`.
   d. In the Password field, specify the password for the Administrator DN. For this scenario, this is `iscpass`.

Note:

- This password, similar to all those in this scenario, are for example purposes only. Do not ever use this, or any other published password, in your own installation or configuration of a product or service.
- Ensure that you enter the correct user ID and password. If you do not, the installation displays a warning message and you must return to the previous page to enter the correct information.

e. Click **Next**.

10. In the LDAP Security Settings page, specify the following information:
    a. In the DN for user and group tree field, specify the fully qualified DN for the top of the LDAP directory tree that holds users and groups objects. For this scenario, we enter `o=myitsoco,c=us`.
    b. In the User naming attribute field, enter the LDAP attribute that maps the login user name to the corresponding LDAP user object. For this scenario, we enter `uid`.
    c. In the Relative location of users field, enter the location of users in the directory. This value is a distinguished name relative to the DN of user and group tree. For this scenario, we enter `cn=users`.
    d. In the User object class field, enter the object class that defines which user objects can be used for authentication. For this scenario, we enter `inetOrgPerson`.
    e. In the Group naming attribute field, enter the LDAP attribute that can map a group short name to the corresponding LDAP group object. For this scenario, we enter `cn`.
    f. In the Relative location of groups field, enter the distinguished name that indicates the location of groups in the directory. This setting is a distinguished name relative to the DN of user and group tree. For this scenario, we enter `cn=groups`.
    g. In the Group object class field, enter the object class that defines which group objects can be used for access control. For this scenario, we enter `groupOfUniqueNames`.
    h. In the Group membership attribute field, enter the LDAP attribute that defines group membership. For this scenario, we enter `uniqueMember`.
    i. Click **Next**.
11. On the User Bind Information for LDAP Server page, use the information from your planning worksheets to specify the following:
    a. In the DN for user and group lookup, enter the fully qualified DN for the server to use to authenticate to the Directory Server during user authentication. This ID is the initial ISC log in ID. For this scenario, we enter `uid=iscadmin,cn=users,o=myitsoco,c=us`.
    b. In the Password field, enter the password for this LDAP DN. For this scenario, we enter `iscpass`.

    Note: This password, similar to all those in this scenario, is for example purposes only. Do not use this password or any other published password in your own installation or configuration of a product or service.

    c. Click **Next**.
12. Review the information about the LTPA Settings page, which describes the LTPA settings needed for the WebSphere Application Server that the ISC and Virtualization Engine console use. Click **Next**.

13. In the LTPA Security Settings page, use the information from your planning worksheets to specify the following information:
    a. In the LTPA password field, enter the password you want to use to create the LTPA key file. For this scenario, we enter `my4password`.

       Note: This password, similar to all those in this scenario, is for example purposes only. Do not use this password or any other published password in your own installation or configuration of a product or service.

    b. In the Confirm password field, re-enter the password.
    c. In the LTPA timeout (seconds) field, enter the number of minutes that you want LTPA tokens that are passed among servers for authentication to be valid. For this scenario, we accept the default value of 120 minutes.
    d. In the SSO domain name field, enter the domain name for your network or enterprise. This domain name designates the domain name of the LTPAToken cookie, and LTPA tokens are accepted by the server only when they come from this domain. For this scenario, we enter `myitsoco.com`. The default value for this field is the domain name for the current installation.

       Note: To ensure that you can use the Virtualization Engine console across your entire network or enterprise, you should enter the broadest possible valid domain name rather than a fully-qualified domain name which might represent only a portion of your network or enterprise.

    e. Click **Next**.
14. Review the text page about the Integrated Solutions Console Configuration page and click **Next**.
15. On the TCP/IP Connection Information for the ISC page, use the information from your planning worksheets to specify the following information:
    a. In the ISC host name field, enter the fully-qualified host name of the system on which you are installing the ISC. Alternatively, you can verify and accept the default host name. For this scenario, we enter `mgmtsrv1.myitsoco.com`.

       Note: Do not enter the system host name only. If you do not enter the fully-qualified host name, the ISC installation will fail.

    b. In the Starting port field, enter the first port number of the nine sequential ports that you want the ISC to use. For this scenario, we accept the default value, which is `8421`.
    c. Click **Next**.
16. On the ISC - LDAP Security page, use the information in your planning worksheets to specify the following information:
    a. In the Portal administrator DN field, enter the fully qualified LDAP DN that you want to use to log in to the ISC and Virtualization Engine console. This is the user ID from the LDAP Security Settings page. Alternatively, you can verify and accept the default DN. For this scenario, we enter `uid=iscadmin,cn=users,o=myitsoco,c=us`.

Note: This DN is given Portal and WebSphere Application Server administrator authority. The WebSphere Application Server runs under this DN and the server cannot start if it cannot use this DN to authenticate to the Directory Server. Record this information because you will need it to stop the WebSphere Application Server instance if you uninstall the Virtualization Engine.

b. In the Short user name field, enter the user ID to use for the Portal administrator DN when accessing the ISC. Or, you can verify and accept the default name. For this scenario, we enter `iscadmin`.

c. In the Portal administrator password field, enter the password for the specified portal administrator DN. Or, you can verify and accept the default name. For this scenario, we enter `iscpass`.

Note: This password, similar to all those in this scenario, is for example purposes only. Do not use this password or any other published password in your own installation or configuration of a product or service.

d. In the Portal administrator group DN field, enter the group DN for an existing LDAP group object whose members are to have Portal administrator authority. Alternatively, you can verify and accept the default DN. For this scenario, we enter `cn=iscadmins,cn=groups,o=myitsoco,c=us`.

e. On the Short group name field, enter the short common name (cn) as the group ID name to use for the specified portal administrator group DN. Alternatively, you can verify and accept the default DN. For this scenario, we enter `iscadmins`.

f. Click **Next**.

17. When the ISC Port Assignment Summary page displays, record the ISC port numbers for future reference.

18. Review the Virtualization Engine console Configuration page and click **Next**.

19. On the Virtualization Engine console – LTPA Key File page, select to create and Export LTPA cryptographic keys to a file, and click **Next**.

Note: These keys are used for security among the various WebSphere Application Servers in the Virtualization Engine. You need to import this key file later in the installation process for use with other WebSphere Application Servers for your Virtualization Engine bridge services. Consequently, when you finish the installation, you need to place this key file on a network file server so that it is readily available for later installations.

20. On the Virtualization Engine console - Export LTPA Key File page, enter the full path and file name for the export key file. You can also click Browse to select a different location. MyITSOCo accepts the default location of /QIBM/ProdData/VE. Record this path and file name so that you can import the file during later installations.

21. On the Virtualization Engine console bridge to IBM Director page, click **Next**.

22. On the Create WebSphere Application Server instance page, specify the following settings:
    a. In the Starting Port field, specify the beginning number for a block of 12 sequential ports for this WebSphere Application Server instance to use. For this scenario, we accept the default value of 11994.
    b. Select **Start instance when system is booted?**
    c. Click **Next**.
23. Review the text on the next page and click **Next**.
24. In the Server Connection Information page, specify the following information:
    a. In the Director Server host name field, specify the fully-qualified host name for the Director Server that you have configured.
    b. In the Director user field, specify the user name for the Director Server.
    c. In the Director login password field, specify the password for the Director Server user identity. For this scenario, we enter `pass`.

    Note: This password, similar to all those in this scenario, is for example purposes only. Do not use this password or any other published password in your own installation or configuration of a product or service.

    d. In the Verify password field, enter the password again.
    e. Select **Connect to Director Server** to validate information field.
25. Review the text on the next page to view which values will be used to create the WebSphere Application Server instance and click **Next**.
26. Review the Virtualization Engine Installation Summary and click **Next**.

    Note: The installation process can take some time. How much time the installation runs depends on your system's processor speed and the size of system memory. It is not uncommon for the installation to take two to four hours.

27. On the Virtualization Engine console page, record the URL that is provided for the Virtualization Engine console interface:

    `http://`*`systemname`*`:892/ibm/console`

    In this URL, *systemname* is the fully qualified system name for this system. For MyITSOCo, Inc., this is `mgmtsrv1.myitsoco.com`. If you selected other Virtualization Engine systems services to install, such as EWLM agent nodes or console bridges, deselect **Launch browser** before you click **Next** to continue the installation of these systems services.
28. Review the information about the EWLM post installation configuration and click **Next**.
29. Read the information about how to launch the configuration wizard for the EWLM Domain Manager and click **Next**.
30. Review the information about how to install managed servers in the next page, record the path and file name for where the managed server images are located, and click **Next**.

31. Review the information in the Installation Completed page, and click **Next**.

    This page displays every systems service that installed successfully. In this situation, MyITSOCo sees the following list:

    – Virtualization Engine base support
    – WebSphere Application Server
    – Integrated Solutions Console
    – EWLM
    – Virtualization Engine console
    – Virtualization Engine console bridge to IBM Director
    – Virtualization Engine managed server installation images

    Note: If you do not need to install any other Virtualization Engine systems services, you can delete the Virtualization Engine media images that you copied or downloaded for the installation.

Step 3: Verifying that you can access and using the Virtualization Engine console

After you complete the Virtualization Engine Installation wizard, you need to verify that the installation was successful by signing on to the ISC and accessing the Virtualization Engine console for the first time. The console is a browser-based management tool that allows you to create and view monitors and other management tasks for various systems.

To launch the console and verify that it performs correctly, complete these steps:

1. Launch a browser session and open the ISC URL:

   `http://`*`your.server.name`*`:`*`isc_port`*`/ibm/console`

   In this URL, *your.server.name* is the fully qualified host name for the ISC installation and *isc_port* is the port for the ISC.

   You must specify the protocol name (`http`) in the URL, because the URL contains a port number. For OS/400 installations, the port number is 2001. For the other platforms, the default port is 8421.

   Do not specify localhost or the IP address for the host name. Doing so prevents you from logging in. For this scenario, we enter `mgmtsvr1.myitsoco.com:892/ibm/console`.

2. When the Virtualization Engine console displays, log in as the ISC administrator. Specify the user ID that you configured during the Virtualization Engine console installation. The default user ID is `iscadmin`. For this scenario, we enter `iscadmin` for the user name and `icspass` for the password.
3. To create new Virtualization Engine console users and groups, go to the Settings tab to display the Settings navigation tree.
4. Click **User and Group Management**.
5. Click the Portal administrator group short name, which by default is **iscadmins**. (You might have to scroll down to find this group.)
6. Click **New User**.
7. Complete the required information and click **OK** to create a new entry in the LDAP server under the Portal administrator group DN.

Note: To view help for the portlet on the page, click the portlet help icon — the question mark (?). The help topic Manage users and groups displays in a separate browser window. Follow the instructions for creating new users and groups and close the help window when you finish. For more information about the Integrated Solutions Console toolbar, click Help from the Integrated Solutions Console toolbar. A separate window is launched and displays a navigation frame for accessing the Console Basics help set.

8. To log out of the console, click **Log off** on the toolbar and return to the Login page.

After you verify that you can access and use the Virtualization Engine console, you need to install bridges on your management sources so that you can use the console to work with them. See the Connecting to your management sources topic in the Virtualization Engine console topic collection for more information.

You have now completed the Virtualization Engine installation for Management Server 1. In the next tasks, you install other Virtualization Engine systems services on other systems in the enterprise. These services allow other systems in the enterprise to interact with the Virtualization Engine console and EWLM on the Management Server.

8.2.4 Task 4: Installing and configuring the Virtualization Engine agents on managed servers

This section lists the steps that you need to perform to install and configure the Virtualization Engine agents on managed servers.

Step 1: Distributing the IBM Director Console install images

Distribute the IBM Director Console install images to the PCs that you plan to use for Director management consoles. You can do this by using FTP with the binary option, by mapping a drive and copying the files to the file system, or by using any other distribution mechanism that you might have in place. The IBM Director Console install image to transfer is dir4.20_console_windows.zip.

Step 2: Installing IBM Director Console on a system running Windows

The IBM Director Console allows you to access and manage system information that the IBM Director Server gathers about systems within the domain. To install the IBM Director Console, complete these steps:

1. Unzip dir4.20_console_windows.zip.
2. Run IBMDirectorConsole4.20.exe to launch the Install Shield Wizard. Follow the instructions in the wizard to complete the installation.

Step 3: Distributing the IBM Director Agent images

Distribute the IBM Director Agent images to the servers that are managed servers in your IBM Director environment. You can do this by using FTP with the binary option, by mapping a drive and copying the files to the file system, or by using any other distribution mechanism that you might have in place. Which IBM Director Agent installation images you need to transfer to a system depends on its operating system image. For this scenario, we transfer the following files:

- dir4.20_agent_ppclinux.tar.gz
- dir4.20_agent_os400.sav
- Q5733VE1110.savf

Step 4: Installing IBM Director Agent to each managed server that is running OS/400

To install IBM Director Agent to each managed server that is running OS/400, follow these steps:

1. Make sure that the following stream files for the Virtualization Engine base code and IBM Director Agent exist in a directory on the managed system (for example /tmp):
   - Virtualization Engine base code: Q5733VE1110.savf
   - IBM Director Agent image: dir4.20_agent_os400.sav
2. On the managed server, copy the installation image stream files into save files using the Copy From Stream File command, **CPYFRMSTMF**. For example:

```
CPYFRMSTMF FROMSTMF('/tmp/Q5733VE1110.savf')
   TOMBR('/qsys.lib/qgpl.lib/ve1base.file') MBROPT(*REPLACE)
CPYFRMSTMF FROMSTMF('/tmp/dir4.20_agent_os400.sav')
   TOMBR('/qsys.lib/qgpl.lib/agent.file') MBROPT(*REPLACE)
```

3. Install the IBM Director Agent using the **RSTLICPGM** command. You must first restore the Virtualization Engine base code, 5733VE1 *BASE, then the Director Agent, 5733VE1, option 39. For example:

```
RSTLICPGM LICPGM(5733VE1) DEV(*SAVF) OPTION(*BASE)
   RSTOBJ(*ALL) LNG(2924) SAVF(QGPL/VE1BASE)
RSTLICPGM LICPGM(5733VE1) DEV(*SAVF) OPTION(39)
   RSTOBJ(*PGM) LNG(2924) SAVF(QGPL/AGENT)
```

4. After you install the IBM Director Agent, you can delete the save file and stream file to save disk space. You can delete the save file using the **DLTF** command. If you moved the stream file to the managed systems, then you can use the **DEL** command to delete it.

Step 5: Installing IBM Director Agent to each managed server that is running Linux

To install IBM Director Agent to each managed server that is running Linux, follow these steps:

1. Decompress the dir4.20_agent_ppclinux.tar.gz file.
2. At the command prompt, specify the name of the installation executable, ./IBMDirectorAgent4.20-1.sh, and press Enter. This executable file launches an unattended installation.

Step 6: Configuring IBM Director on systems that are running OS/400

To configure IBM Director on systems that are running OS/400, follow these steps:

1. Enable encryption for IBM Director. Complete the following steps on each management server and managed server running OS/400:
   a. Using an ASCII text editor if you have a mapped drive or the **EDTF** command in the character-based interface, open the /QIBM/ProdData/Java400/jdk13/lib/security/java.security file.
   b. Uncomment the following line by removing the hash (#) at the beginning of the line: `security.provider.3=com.ibm.crypto.provider.IBMJCE`.
   c. Run the following Qshell command to enable encryption:

```
/qibm/userdata/director/bin/cfgsecurity
```

2. Start IBM Director. From the character-based interface, enter the following command:

   ```
   STRTCPSVR SERVER(*DIRECTOR)
   ```

3. Authorize OS/400 users to IBM Director functions by registering them in a function usage group: IBM Director Super Administrators or IBM Director Administrators. To authorize users to these functions, you must have *SECADM authority. Complete the following steps to authorize users:

   a. In iSeries Navigator, right-click the server and click **Application Administration**.

   b. On the Application Administration dialog, go to the Host Applications tab.

   c. Expand IBM Director for iSeries.

   d. Select the function group to which you want to add users and click **Customize**. Complete the instructions on the dialog to grant authority.

Step 7: Configuring IBM Director on systems that are running Linux in the IBM Director domain

To configure IBM Director on systems that are running Linux in the IBM Director domain, do the following:

1. Enable encryption for IBM Director. Enter the following command:

   ```
   /opt/ibm/director/bin/cfgsecurity
   ```

2. Start IBM Director. Enter the following command:

   ```
   /opt/ibm/director/bin/twgstart
   ```

Step 8: Starting IBM Director Console

To start IBM Director Console, do the following:

1. To start IBM Director Console on a system running Windows, click **Start** → **Programs** → **IBM Director** → **IBM Director Console**. Then, do the following:

   a. In the IBM Director Server field, enter the name of the management server.
   b. In the User ID field, enter a valid user ID for IBM Director.
   c. In the Password field, enter the password that corresponds to the user ID.
   d. Click **OK**.

2. Complete the Event Action Plan wizard that starts every time you log in to IBM Director Console until you take action.

3. Discover systems in your IBM Director environment that have IBM Director Agent installed. Click **Tasks** → **Discover Systems** and select the type of systems that you want to discover.

8.2.5 Task 5: Configuring EWLM Domain Manager

To configure the domain manager, run the following command on the same server where the domain manager is installed:

```
QWLM/STRWLM CFGID(myDomainManager) MODE(*DMNMGR) DMNMGRPORT(port) ADMPORT(port)
DOMAIN('DomainName')
```

Note: There are other, optional parameters that you can set with this command. A list of all the possible parameters that are required and under what circumstances you should use them is available in the **STRWLM** topic.

8.2.6 Task 6: Installing and configuring EWLM Managed Servers

To install and configure EWLM Managed Servers on each managed server, complete the steps in this section.

Step 1: Distributing managed server images to each server that is part of the EWLM domain

Distribute the managed server images to each of the servers that will be part of an EWLM management domain. MyITSOCo distributed the images by using FTP with the binary option. You can also do this by mapping a drive and copying the files to the file system or by using any other distribution mechanism that you might have in place.

Step 2: Installing the EWLM Managed Servers (in interactive mode) on each managed server that is running OS/400

To install the managed server, complete these steps:

1. From the directory on the system running Windows where you copied the EWLMOS400MS.jar and Q5733VE1110.savf files, run this command:

   ```
   java -jar EWLMOS400MS.jar -os400 target-system-name user-ID password
   ```

   In this command, *target-system-name* is the host name of the target OS/400 system, and *user-ID* and *password* are the user ID and password for accessing the target OS/400 system. Ensure that the user profile under which you sign on has the proper authorities for installing programs.
2. Complete the Install wizard to install the managed server on the target system.
3. Repeat these steps for each managed server.

Step 3: Configuring OS/400 as a managed server

To configure OS/400 as a managed server, follow these steps:

1. Run this command to configure a managed server on OS/400:

   ```
   QWLM/STRWLM CFGID(myManagedServer) MODE(*MGDSVR) DMNMGRPORT(port)
   DMNMGRNAME('DomainName')
   ```

   Note: There are other, optional parameters that you can set with this command. A list of all the possible parameters that are required and under what circumstances you should use them is available in the **STRWLM** topic.

2. Repeat these steps on each managed server.

Step 4: Configuring the EWLM Control Center on Management Server 1

The EWLM Control Center is a Web-based user interface and is hosted on the server that you identified as the domain manager. You must use the createCC.sh configuration script to create a WebSphere Application Server instance, under which the EWLM Control Center will run. After you create the EWLM Control Center on OS/400, you must assign users to EWLM Control Center roles and then start it.

Step 5: Starting EWLM

EWLM on OS/400 is started automatically when you run the **STRWLM** command to set up your domain manager, single-system EWLM Managed Servers or firewall brokers.

8.2.7 Task 7: Beginning the planning process for implementing systems provisioning

Systems provisioning allows you to designate a pool of resources that are shared between different workloads. You can then dynamically assign the resources of these servers, along with their major operating systems and existing virtualization technologies, such as hypervisor, LPAR, virtual machine, and so on, to workloads that need capacity. Little or no human interaction is needed. Systems provisioning is a sequence of activities that must occur in a specific order. The activities are defined in a workflow, which is run by a workflow engine.

Now that MyITSOCo, Inc. has installed, configured, and verified the Virtualization Engine environment, they begin the process of planning how to add systems provisioning to the Virtualization Engine environment. They use the systems provisioning planning guide to plan their installation and configuration of Tivoli Provisioning Manager on their managed-through server, which is named Tivoli Provisioning Manager Server.

This planning guide describes systems provisioning in the context of the Virtualization Engine. It contains planning information for those who want to install systems provisioning. Tasks described in the planning guide include:

- Determining a possible topology for deployment of provisioning.
- Determining the hardware and software requirements for provisioning.
- Using scenarios to determine a deployment plan.

See the Systems provisioning topic collection to learn more about systems provisioning, including how to install and use Tivoli Provisioning Manager.

9

Installing Virtual Machine Manager and VMware on managed servers

In the context of the IBM Virtualization Engine the main focus resides on the ability to virtualize systems and enterprises. In this chapter we focus on an already known way of virtualization that was brought up by companies such as VMware with their products in virtualizing Intel and x86 type architectures on a large scale. Nowadays VMware is not the only vendor of virtualizing these architectures but the only two that IBM has started to integrate into the Virtualization Engine platform: VMware with VirtualCenter can play a very complementary role with the VMware ESX Server and Microsoft Virtual Server.

The focus on this chapter is how to integrate virtualization software into the Virtualization Engine suite considering using an VMware example.

© Copyright IBM Corp. 2005. All rights reserved.

9.1 Virtual Machine Manager

The Virtual Machine Manager (VMM) is a *software enhancement* to IBM Director that integrates the virtualization software ESX Server through VirtualCenter from VMware and Virtual Server from Microsoft into the IBM Director managed environment. In this redbook, we have chosen the VMware setup as an example.

IBM Director is one of the key components of the set of systems management features that are available on the xSeries and BladeCenter platforms. This tool is aimed at the management of a system that spans from the hardware to critical functions of the operating system. It provides:

- A single consistent interface for managing an entire server farm, keeping an up-to-date inventory of the current infrastructure.
- An engine that can respond and react when a problem occurs on the infrastructure.

VMM extends the IBM Director to include *virtualization awareness*. Specifically, VMM supports both VMware ESX server deployments and Microsoft Virtual Server deployments. It also provides a detailed view of physical platforms correlated to the virtual environment that each of those supports.

In VMM Release 1.0, the system is able to detect software and hardware alerts (such as a PFA alert or another non-critical hardware issue) and to migrate one or more virtual machines off that server onto another server that is operating normally and has sufficient resources to support the additional workload.

The following sections introduce the installation of all software components in the IBM Director environment using VMware ESX server to run virtual machines with Windows guest operating systems.

9.2 The topology

This chapter is a continuation of Chapter 7, “Installing the Virtualization Engine console, IBM Director, TPM, and EWLM on Blade with Windows” on page 161. The enhancement that we introduce here is the use of VMware running on a dedicated server. Therefore, we use all the installed and running components of the Chapter 7, “Installing the Virtualization Engine console, IBM Director, TPM, and EWLM on Blade with Windows” on page 161 installation. This chapter only covers the additional hardware and software components that were added after building and configuring the Virtualization Engine for Windows environment. If you need more information about how this Windows setup is built, see the referenced chapter.

The following are the enhancements that we make to the existing Windows setup. The changes are listed per server to make it easier to get an overview of how the new components interrelate to each other and to the existing server infrastructure:

- VEWIN21 - IBM Director Server

  The VMM server extension allows the user to manage virtual resources. The managing server that runs the IBM Director server function is extended through the VMM Server component. Through this enhancement, the IBM Director is able to understand the communication with the VMM components in the IBM Director’s domain.

- VEWIN21 - IBM Director Console

  The VMM console extension enables the IBM Director Console to perform various management tasks on the virtual resources being managed. In our topology, the IBM

Director console resides on that server and therefore the VMM Console package is installed as well as the VMware VirtualCenter client software.

- VEWIN22

  This system has the VMware VirtualCenter server and the VMware VirtualCenter Web Service installed. The IBM Director 4.20 Agent is installed along with the VMM Agent extension. The Director Agent and VMM Agent extension serve as the management interface between the VMware VirtualCenter Web Service and the IBM Director Server with the VMM Server extension providing management of the virtual resources on the VMware hosts.

- MSNESX1

  VMware ESX Server is installed on this system. The IBM Director Agent for Linux is installed on this system so that IBM Director Server can manage the ESX Server. This Linux Agent provides Director Server with discovery capabilities and basic management functionality but is not virtual resource aware. The virtual resources on this ESX host will be managed by VMware VirtualCenter, the IBM Director Agent and the VMM Agent Extension this provides a single point of management for both physical and virtual resources for this ESX host. This management point is exposed at the Director Console where the VMM extension is installed.

- VM1 and VM2

  Are the two virtual machines that are defined on the ESX Server MSNESX1. These virtual machines are managed at the same management point from which the ESX host is managed. The definition of a virtual machine within VMware ESX Server encompass all physical parameters a hardware could have, such as the number of processors, the amount of memory, the network cards, and so forth This parameter represents the runtime boundaries of the virtual hardware.

- MSNVMG1

  It is a VMware guest operating system with a version of Windows Server 2003 Standard Edition installed on. MSNVMG1 is running in a virtual machine VM1 defined on the VMware ESX Server. Additionally, the IBM Director Agent is installed on this guest operating system instance to allow IBM Director direct management of this virtual server.

- MSNVMG2

  This guest operating system is a clone of MSNVMG1. Note: take caution when cloning virtual machines with director agents installed; there might be some identifier conflict exposed by cloning that causes some problems with managing via Director.

Figure 9-1 illustrates a detailed overview of the topology.

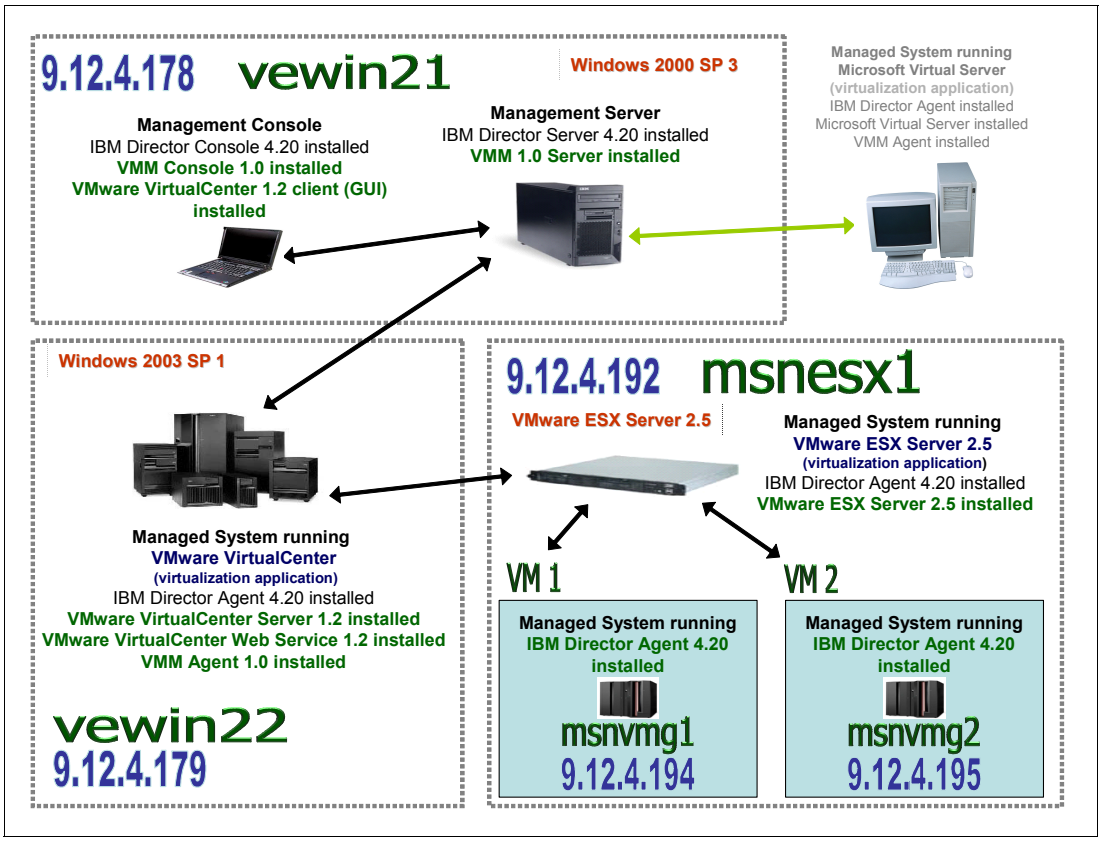

Figure 9-1 System context for the VMM setup

Hardware and software considerations

From a topology standpoint, some consideration needs to be well understood:

- We have to add a VMware ESX Server to the Windows topology. The ESX Server only supports SCSI hard discs. And therefore the software could only be installed if SCSI discs are available in the server.
- Another constraint of VMware ESX Server, in a pure Ethernet environment is you need to have two active Ethernet cards into the server. The reason is performance: the communication between the managing hosts (such as a VMware VirtualCenter and IBM Director) and the traffic on the guest operating system must be on separate network cards.

  We decided for the ITSO setup to run with one single network card because we did not have the need of most optimal network communication speed.
- The VMware ESX Server is a software package that installs on a Linux based operation environment. That requires a single server that only runs ESX Server from VMware. At the time of writing the only VMware product that is supported from IBM with VMM 1.0 is the combination of IBM Director 4.2 plus VMware VirtualCenter 1.2 plus ESX Server 2.5.

  The reason why we have chosen VMware ESX Server is that the other VMware products, GSX Server or VMware Workstation, are not supported by VMM right now.

9.3 The VMM installation process

The following installation of the VMM suite builds on the installation that has been documented in the Chapter 7, “Installing the Virtualization Engine console, IBM Director, TPM, and EWLM on Blade with Windows” on page 161. Therefore, this section describes only the additional steps to implement the VMM scenario.

The following lists the hardware and software components that are described in a sequence:

1. Setting up the VMM environment:
   a. Installing VMM Server on IBM Director Server node
   b. Installing VMM Console on IBM Director Console node
   c. Installing VMM Agent on VMware VirtualCenter Server node
2. Deploying of the additional IBM Director Agent
   a. Installing the IBM Director Agent on VMware ESX Server
   b. Installing the IBM Director Agent on guest operating systems running on VM Ware ESX Server

Note: Appendix B, “VMware installation” on page 367 describes the full installation of the VMWare set processes.

9.3.1 Setting up the VMM environment

VMM enables VMware Virtual Center and Microsoft Virtual Server to be managed by some means through an IBM Director. The scope of managing features is dependent on the virtual vendor (VMware or Microsoft). For more details, see the IBM VMM User’s guide. For the latest information and documentation, refer to the following:

http://www-307.ibm.com/pc/support/site.wss/document.do?lndocid=MIGR-56914

To enable VMM in an IBM Director environment you have to install three different modules:

VMM Server extends the IBM Director objects adding VMM object attributes and context menu items to the virtual objects as they are identified and provides the interface for communicating with the VMM Agent. VMM Server has to be installed on the same server where the IBM Director Server is installed. In our topology this is server vewin21.

VMM Console provides a mechanism to launch vendor software for accessing the third party management interfaces. It also provides dialogs to log in to the VMware Web Service interface and add host to VMware VirtualCenter. Thus the IBM Director Console is able to display the new VMM objects and could also build associations that are specific to VMM objects. Also all enhancements in managing and querying the VMM objects are build into the IBM Director Console through the installation. VMM Console has to be installed on the same server where the IBM Director Console is installed. In our topology this is server vewin21.

VMM Agent is the interface to the virtual machine managing product (in our setup VMware). The VMM agent is the counterpart of the VMM Server installed on IBM Director Server. A request is sent from the IBM Director Server to the corresponding VMM agent that uses the VMware Web Service interface to post this request onto the VMware VirtualCenter Server. The result of that request is sent back to the requestor. VMM Agent has to be installed

on the same server where VMware VirtualCenter Server is installed. In our topology this is server vewin22.

Note: VMM 1.0 is an enhancement to the Windows IBM Director and the virtualization applications from VMware and Microsoft. To get the list of supported Windows operating system versions, consult the corresponding product manuals of IBM Director 4.2, VMware VirtualCenter and Microsoft Virtual Server. For the latest information and documentation, refer to the following:

http://www-307.ibm.com/pc/support/site.wss/document.do?lndocid=MIGR-56914

The VMM installer can be downloaded for free. For the latest information and documentation, refer to the following:

http://www-307.ibm.com/pc/support/site.wss/document.do?lndocid=MIGR-56914

The installation image is a Windows executable code that could be started through the Microsoft Explorer or the Start → Run command.

Dependent on the already installed components on the server, the Installer enables only those components that are able to be installed. This means that if the Installer detects that no IBM Director Server is installed, the option to install VMM Server will not be available.

The following sections describe how to install the VMM components on a given server infrastructure.

Setting up VMM Server on IBM Director Server node

To get the VMM enhancements on an IBM Director managed system available you have to install the VMM Server portion on that IBM Director Server.

Planning VMM Server installation

Before you install, check the following:

- The installation of the VMM Server does not need any parameter beside the information of where to store the VMM Server application but it has to stop the IBM Director Service on that node. Therefore if you plan to install VMM Server be sure that you could stop the IBM Director Service for a couple of seconds.
- The Windows user that installs VMM Server has to be a member of the **Administrators** group.

Installing VMM Server

The VMM Server installation process guides you through different screens. Only the major steps that are important to define the parameters that configure the server instance are listed below; for instance one screen such as the End User Licence Agreement window was left out. For more information about installing VMM Server, go to the IBM Virtual Machine Manager 1.0 Installation documentation. For the latest information and documentation, refer to the following:

http://level3.kirkland.ibm.com/VMWare/

To move though the setup, select **Next**.

These are the major configuration steps:

1. IBM Virtual Machine Manager Welcome window: Click **Next**.
2. Setup Type window: For installing VMM Server, you have to select **Install VMM Server and Console** and click **Next**.

3. Destination Folder window: Choose the given recommendation and click **Next**.
4. Ready to Install the Program: Click **Install**.
5. Question window: The installation of VMM Server needs to stop the IBM Director service on that server. If you could not allow the installation program to do so click **No** and install VMM Server later. Otherwise, click **Yes** and the installation is performed.
6. IBM Virtual Machine Manager Completion window: The installation is complete. Click **Finish** to exit the installer.

The time needed for copying the files, stopping the IBM Director Service, and performing the configuration tasks on the Windows operating system takes between one to two minutes, depending on the speed of the processor and hard disks in the system.

Verifying VMM Server installation

To check whether the installation of VMM Serer features succeeded successfully, you can do the following:

- From a program installation perspective:

  Through the Windows system **Start** → **Control Panel** → **Add or Remove Programs** there is one row that has the label **IBM Virtual Machine Manager**. Click the row of VMM and then select **Click here for support information**. The Support Info window opens and shows you the product information of VMM with some other useful information, as shown in Figure 9-2.

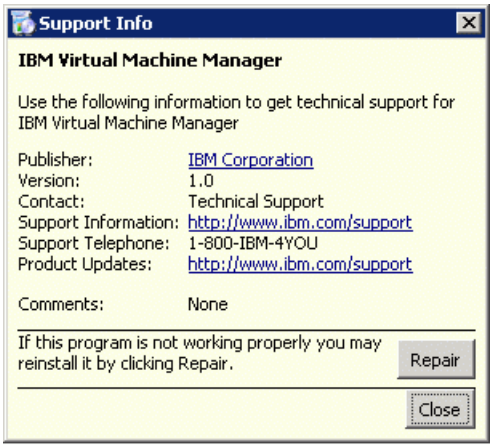
Support Info
IBM Virtual Machine Manager
Use the following information to get technical support for IBM Virtual Machine Manager
Publisher: IBM Corporation
Version: 1.0
Contact: Technical Support
Support Information: http://www.ibm.com/support
Support Telephone: 1-800-IBM-4YOU
Product Updates: http://www.ibm.com/support
Comments: None
If this program is not working properly you may reinstall it by clicking Repair.
Repair
Close

Figure 9-2 IBM VMM Support Information

- From an IBM Director Server perspective:

  Open IBM Director Console, go to the menu and choose **Help** → **Product Information**. The IBM Director Product Information window opens, and through the installation of VMM Server, a new Extension appears in the list Virtual Machine Manager from IBM in Version 1.0.

The next step is to install the VMM Console on the IBM Director Console computer to get VMM objects displayed on the console and to start interacting with them from an management perspective.

Note: In our environment (the IBM Director Server and the IBM Director Console are on one server that is named vewin21), the VMM installer has already installed the VMM Server part and the VMM Console part in one single installation process. If you have the same environment, you can skip to “Verifying the VMM Console installation” on page 233.

Setting up the VMM Console on the IBM Director Console node

The VMM Console Extension must be installed on the IBM Director Console to provide the interface for managing the virtual resources.

Planning the VMM Console installation

Before you start the installation, check the following:

- The installation of the VMM Console does not need any parameter beside the information of where to store the VMM Console application.
- The Windows user that installs VMM Console has to be member of the **Administrators** group.

Installing the VMM Console

The VMM Console installation guides you through different screens. Only the major steps that are important to define the parameters that configure the server instance are listed below. For example one window, such as the End User Licence Agreement window, was left out. For more information about installing VMM Server, go to the IBM Virtual Machine Manager 1.0 Installation documentation:

http://www-307.ibm.com/pc/support/site.wss/document.do?lndocid=MIGR-56914)

To move through the process, click **Next**.

These are the major configuration steps:

1. IBM Virtual Machine Manager Welcome window: Click **Next** .
2. Setup Type window: For installing VMM Console, you have to select **Install VMM Console only** and click **Next** .
3. Destination Folder window: Chose the given recommendation and click **Next** .
4. Ready to Install the Program: Click **Install**.
5. The installation is performed.
6. IBM Virtual Machine Manager Completion window. The installation is complete. Click **Finish** to exit the installer.

The time needed for copying the files and configurations tasks on the Windows operating system takes between one to two minutes, depending on the speed of the process and hard disks in the system.

Verifying the VMM Console installation

Tip: After the installation of the VMM Console on the server, it is necessary to restart the IBM Director Server service to make the new VMM objects visible in the IBM Director Console.

After the service is restarted, new VMM objects show up in the IBM Director Console as follows:

- In Groups: VMM Systems is a new group (filter) that displays only those systems in the Group Contents view that are valid coordinators.

Note: A coordinator is a system on which all the following components are installed:

- IBM Director Agent
- VMM Agent
- Virtualization application: VMware VirtualCenter Server with Web Service

- In Group Contents: All systems are displayed that are valid coordinators. You must have first installed VMM Agents to see coordinator objects.
- In Tasks: Virtual Machine Manager with corresponding subtasks Help and Start Vendor Software.

Note: The IBM Director Console has the ability to order the objects in the Group Contents column in different ways, depending on the associations you have chosen through the IBM Director Console menu. The installation of the VMM Console software adds a new association VMM Systems Membership if you choose that the VMM objects are organized in a tree structure (depending on the logical structure of the objects).

With this step, the IBM Director installation on the server vewin21 is complete. To get access to VMM object, you now install the VMM Agent on the server that has access to all the virtual machines running in the infrastructure. This is the VMware VirtualCenter server on vewin22.

Setting up the VMM Agent on VMware VirtualCenter Server node

Now, you establish the communication and exchange the information concerning the virtual machines within the organizational structure. The next step is to install the VMM Agent portion on a computer that runs the managing interface of the virtualization application. This establishes the communication with the IBM Director.

Planning the VMM Agent installation

The installation of the VMM Agent does not need any parameter in addition to the information of where to store the VMM Agent application. However, it has to stop the IBM Director Service on that node. Therefore, if you plan to install VMM Agent be sure that you could stop the IBM Director Service for a couple of seconds.

The Windows user that installs VMM Agent has to be a member of the Administrators group.

Installing the VMM Agent

The VMM Agent installation guides you through different steps. Only the major steps that are important to define the parameters that configure the server instance are listed in this section. For example one window, such as the End User Licence Agreement, was left out. For more information about installing VMM Agent go to the IBM Virtual Machine Manager 1.0 Installation documentation.

To move through the process, click **Next**.

These are the major configuration steps:

1. IBM Virtual Machine Manager Welcome window: click **Next** .
2. Setup Type window: To install the VMM Agent you have to select **Install VMM Subagent for VMware VirtualCenter** and click **Next** .
3. Destination Folder window: Choose the given recommendation and click **Next** .
4. Register VMware Certificate with IBM Director window: Choose the given recommendation and click **Next** .

   Tip: The VMM Agent has to use the certificate authority of VMware to establish a secured connection. In this step the VMM Agent installer needs to know in which directory the VMware's root.perm file resides. The suggested directory in the installer points to this location.

5. Ready to Install the Program: Click **Install**.
6. Question window: The installation of VMM Agent needs to stop the IBM Director Agent service on that server. If you could not allow the installation program to do so click **No** and install VMM Agent later. Otherwise, click **Yes** and the installation will be performed.
7. A command prompt window opens and asks if you trust the certificate from VMware in our context. Enter `yes` if you can and press Enter.
8. IBM Virtual Machine Manager Completion window: The installation is complete. Click **Finish** to exit the installer.

The time needed for copying the files, stopping the IBM Director Service, and performing the configuration tasks on the Windows operating system takes between one to two minutes, depending on the speed of the processor and hard disks in the system.

Verifying the VMM Agent installation

Now, all VMM components are installed on the different servers. One way to verify the installation is to look into the system and to list the installed software products. This process is described in "Verifying VMM Server installation" on page 231.

Note: Activate the Association VMM Systems Membership as described in "Verifying the VMM Console installation" on page 233 for convenient display of the VMM objects.

To explore the VMM managed domain start the IBM Director Console on the IBM Director Server (vewin21), do the following:

1. In the column Groups, click the object **VMM Systems**.
2. In the column Group Contents, one system shows up (in our case VEWIN22). The representation of this system is made of two icons:
   - An IBM Director icon to identify servers
   - The identifier that this object is a VMM object (the icon shows a little lock that identifies that the object is locked and the first step to open the object for communication is to enter the credentials for that server).
3. The first action on this VMM object VEWIN22 which represents the VMware VirtualCenter Server is to enter the credentials. This is the first step to open the door for all further management tasks. Therefore, right-click **VEWIN22** and chose in the context menu **Coordinator Management** → **Enter Credentials**.

4. The IBM Director Console opens a new window in which the following values has to be typed in:

   User ID: `vmwadmin`
   Password: `password`
   Port: `8443`

5. Click **OK** and a new widow opens saying **Credentials accepted** as shown in Figure 9-3.

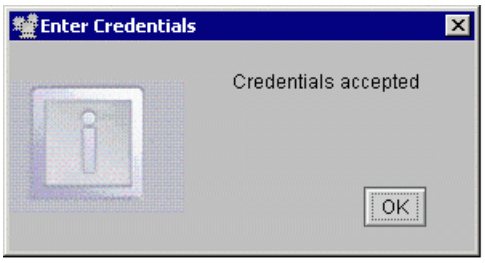

Figure 9-3 Acceptance of the credentials for the VMware VirtualCenter

6. Now you have access to the VMware VirtualCenter objects that are represented as an VMM object in the IBM Director Console. Select the **VEWIN22** object in the Groups Contents column and press the left arrow key on your keyboard. Every time you press the left arrow key, you browse one dimension deeper onto that VEWIN22 server topology (see Figure 9-4).

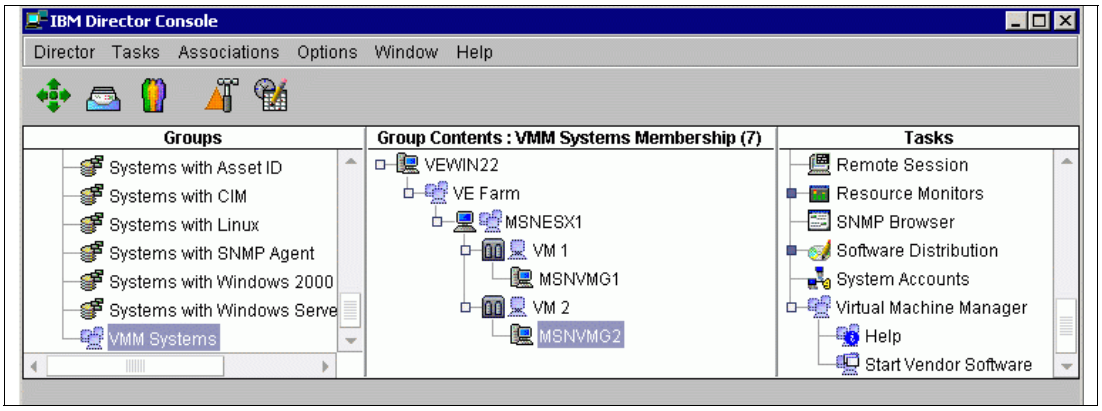

Figure 9-4 All VMM objects on one VMware VirtualCenter server

In this section, only a few options of the IBM Director Console using the VMM enhancements on VMware VirtualCenter environments are shown. For more information in using and configuration go to the IBM Virtual Machine Manager 1.0 Installation and User's Guide. For the latest information and documentation, refer to the following:

`http://level3.kirkland.ibm.com/VMWare/`

To also get the IBM Director capabilities of managing the ESX Server and the virtual machines running on the ESX Server we have to install the appropriate agents. The next section will describe how to.

9.3.2 Deploying an additional IBM Director Agent

To get new servers into the management domain of IBM Director on each of the systems an IBM Director Agent has to be installed. For the two new types of VMware systems:

- The VMware ESX Server
- The virtual machines running on that VMware ESX Server

The following sections describe the installation of the IBM Director Agent.

Setting up the IBM Director Agent on VMware ESX Server

In this section, we discuss the installation of the IBM Director Agent on Linux. For more information how to install IBM Director Agents on Linux consult the IBM Director documentation.

Planning the IBM Director Agent installation on VMware ESX Server

The installation of the IBM Director Agent does not need any parameter. To install the IBM Director Agent for Linux on the VMware ESX Server you need to have the software on a server which could be accessed through ftp from the VMware ESX Server or you have the software on a CD to use the cd-rom drive to access the software.

You need root access to the VMware ESX Server.

Installing the IBM Director Agent on VMware ESX Server

The ESX Server is based on a Linux system. Therefore we need the IBM Director Agents for Linux. The installation of this agent has to take place on the VMware ESX Server physical terminal. The following steps will perform that installation:

1. Log into the VMware ESX Server **MSNVMW1** using the `root` user:

   Login: `root`
   Password: `password`

2. Copy the install images for IBM Director Agent for Linux on this server by using either ftp or copy the file from the CD-ROM in the cd-rom drive to the /tmp directory. The name of the file we used in our setup was dir4.20_agent_linux.tar.gz.
3. Create in the /tmp directory a new directory with the command **`mkdir inst`**.
4. Change directory with the **`cd inst`** command.
5. GNU uncompress the file with the **`gunzip ../dir4.20_agent_linux.tar.gz`** command.
6. Untar the file with the **`tar -xf ../dir4.20_agent_linux.tar`** command.
7. Start the installation in the current shell with the following:

   ```
   ./IBMDirectorAgent4.20-1.sh
   ```

The installation program installs the IBM Director Agent and all needed certificates, as shown in Example 9-1.

Example 9-1 Installation of IBM Director Agent on VMware ESX Server

```
./IBMDirectorAgent4.20-1.sh self-extracting installation program... Please wait...

Attempting to install ITDAgent-4.20-1.i386.rpm
Preparing...                ##################################################
ITDAgent                    ##################################################
To start the IBM Director Agent manually, run /opt/IBM/director/bin/twgstart
Attempting to install DirAgent-4.20-1.i386.rpm
Preparing...                ##################################################
DirAgent                    ##################################################
Compiling mof files...
Creating SSL certificate and private key
Starting Pegasus CIMOM

Compiling sensor mofs...
Either FRU is not supported or this system does not have access to FRU file.
Please run man getfru for help.
GETFRU completed

Installation of selected components is successful.
To start the IBM Director Agent manually, run /opt/IBM/director/bin/twgstart
[root@msnvmw1 inst]#
```

8. The installation of the IBM Director Agent was successful. You now reboot the system with the **reboot** command. Be sure that the VMware ESX Server has stopped all virtual machines and that no user is using this server before you reboot.

The next step is to verify whether the installation was successful and whether the IBM Director Agent is working properly on the VMware ESX Server.

Verifying IBM Director Agent installation on VMware ESX Server

Every time the VMware ESX Server starts, you can observe the progress and status of the boot sequence. If everything is starting as expected, the prompt of the steps initializing a daemon or a function is reported with `[OK]` at the right end of the line. If there is an error, it could display `[FAILED]` and additional error messages occur.

After having installed the IBM Director Agent on a VMware ESX Server one additional status line is displayed on the boot sequence: `Starting IBM Director Agent.` This status indicates that the service has installed successfully and it has loaded without an error.

The next verification step is to start the IBM Director Console on the VEWIN21 server. In the Group Contents list, the VMware ESX Server shows up with the name that is defined in the VMware VirtualCenter in the hierarchy structure (for example, VEWIN22 → VE Farm → msnvmw1.itso.ibm.com).

Tip: If the ESX Server does not show up in the list of groups VMM Systems:

1. Group Contents, you have to go to the IBM Director Console menu and choose **Options** → **Discovery Preferences**.
2. Click **IBM Director Systems** → **System Discovery (IP)**.
3. Select **Add** in the container Address Entries.
4. Chose **Unicast Range** and click **Next**.
5. Enter the Add Unicast Address for the ESX Server IP Address w.x.y.z and click **OK**.

   If you want to add multiple hosts within a IP address range choose the **Unicast Range** choice.

Now, you can start the discovery process from the IBM Director Console again, and the ESX Server shows up.

The access to this new VMM object is not granted so far and this is displayed through a lock at the server. The IBM Director user has to unlock the VMM object through entering the credentials for this VMware ESX Server. To get now access to the VMware ESX Server host `msnvmw1.itso.ibm.com`, you have to enter the credentials for this system. Double-click the icon with the lock symbol and a Request Access to Systems window opens asking you for the User ID `root` and Password `password`. The account needed here is the root account on the ESX Server that we have defined through the VMware Installation.

After pressing Enter you get a message with the result `Access was successfully granted`. The lock symbol in the IBM Director Console disappears.

Because the first VMware ESX Server is added into the managed IBM Director network, an additional icon shows up in the Group column Logical Platforms.

The VMware ESX Server is IBM Director enabled. To integrate the VMware virtual machines into the IBM Director domain, you we have to install IBM Director Agents on each virtual machine.

Installing the IBM Director Agent

The installation of the IBM Director Agents on a Windows 2003 systems has been performed in the Virtualization Engine installation for Windows. Therefore, we refer to this installation to Chapter 7, “Installing the Virtualization Engine console, IBM Director, TPM, and EWLM on Blade with Windows” on page 161

Verifying the IBM Director Agent

After the IBM Director Agents were installed on the VMware guest operating systems the new VMM objects are accessible through the IBM Director Console.

Fully expand the tree structure in the Group Contents column to see the available virtual machines. In our setup, the VMware ESX Server MSNVMW1 has two virtual machines, one named Microsoft Windows Server 2003, Standard Edition that is running a guest operating system with IBM Director Agent named host MSNVMW-1 and the other named Windows 2003 Server Standard Edition that is running a guest operating system with IBM Director Agent named host MSNVMW-2.

9.4 Uninstalling VMM components

The following sections describe how to uninstall VMM on the different server types, management servers, and managed servers.

Note: Every removal of a VMM component needs the IBM Director Service to stop for a short time. This stop of the service has to be planned; the remove procedure will stop and restart the service on its own.

9.4.1 Uninstalling VMM on management servers

In our scenario, the management server is the server named VEWIN21 on which the VMM component Console 1.0 and Server 1.0 are installed on. What had been said for the installation is also true for the de-installation: the VMM Console will be removed in the same step as the VMM Server 1.0.

These are the major steps to remove the VMM Server and Console software:

1. On a Windows system go to **Start** → **Control Panel** → **Add or Remove Programs** and search a row with labeled IBM Virtual Machine Manager.
2. Click **Remove** to start removing the VMM software component on VEWIN21.
3. A window opens with the label Remove Programs. Chose **Yes** to remove the software.

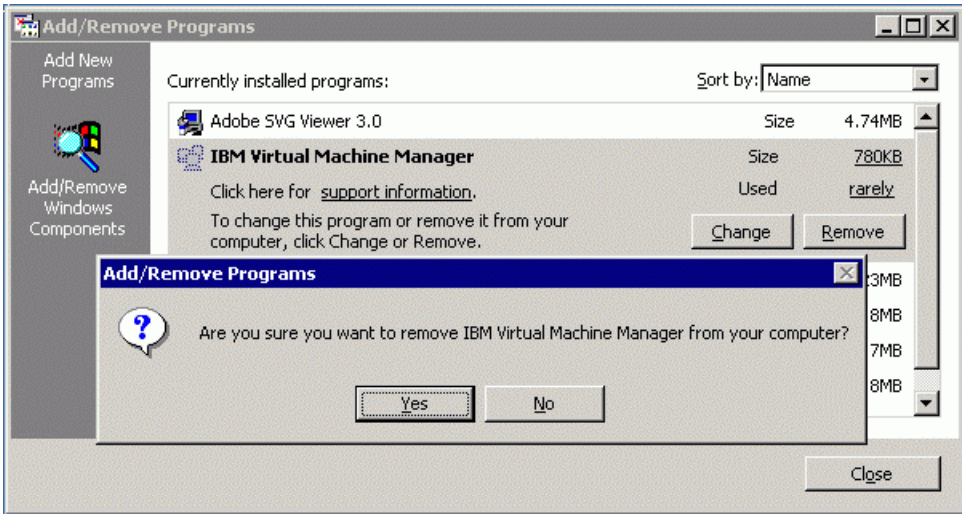

Figure 9-5 Remove VMM Server and Console

4. The remover ask as a precaution if the system can perform the action of stopping the IBM Director Service. If so, answer with **Yes**.
5. All VMM software components are about to be deleted. In the meantime, a dialog box displays that explains that it is about to remove VMM.
6. After the complete removal of the software, the Virtual Machine Manager remover asks if the system could be restarted immediately, because the configuration changes could take effect after restart. If so, answer with **Yes**.

With this last step all VMM Console and Server components are removed from the system VEWIN21.

9.4.2 Uninstalling VMM on managed servers

In our scenario the managed server is the server named VEWIN22 on which the VMM component Agent 1.0 is installed on.

These are the major steps to remove the VMM Server and Console software:

1. On a Windows system go to **Start** → **Control Panel** → **Add or Remove Programs** and search a row with labeled IBM Virtual Machine Manager.
2. Click **Remove** to start removing the VMM software component on VEWIN22.
3. A window opens with the label Remove Programs. Choose **Yes** to remove the software.
4. The remover asks as a precaution if the system can perform the action of stopping the IBM Director Service. If so, answer with **Yes**.
5. All VMM software components are about to be deleted. In the meantime, a dialog box displays that explains that it is about to remove VMM.

   Remark: In our tests, the InstallShield Wizard issued an error message that it could not delete the file IBM Virtual Machine Manager.msi. This is a single file that resides on the system. To delete this orphan file, use the Window system search capabilities and search for the file name. Then, use the delete function of the operating system.

6. After the complete removal of the software, the Virtual Machine Manager remover ask if the system could be restarted immediately, because the configuration changes could take effect after restart. If so, answer with **Yes**.

With this last step the VMM Agent component is removed from system VEWIN22.

Part 3

Using the Virtualization Engine Scenarios

An installation process is not an end in itself. Products are installed to add new functionality, to promote a new service, to solve problems, and so forth. Thus, a product is installed to help the business, not for the sake of the product's wonderful technology.

The Virtualization Engine aims to help systems management in a heterogeneous environment. When the components are installed and configured, the management of non-homogeneous systems is easier, focusing more on business issues than on environment issues.

The chapters in this part describe a few scenarios that demonstrate the value of the Virtualization Engine and include the following:

- Chapter 10, "How to use TPM and EWLM to provision a resource" on page 243
- Chapter 11, "How to monitor an application in a heterogeneous environment using the VE console and EWLM" on page 261
- Chapter 12, "How to fix hardware constraints using TPM and IBM Director" on page 291
- Chapter 13, "How to optimize an xSeries IT environment using IBM Director and Virtual Machine Manager" on page 309

Although we might have considered other scenarios, we chose these scenarios because they include common situations that have been reported by customers.

Important: All these scenarios apply regardless of the platform that hosts the management servers. The functions are independent of the management server platform.

© Copyright IBM Corp. 2005. All rights reserved.

10

How to use TPM and EWLM to provision a resource

In this chapter, we describe how to use Tivoli Provisioning Manager and EWLM to provision a resource when an application does not achieve its goal. It includes the following topics:

- About this scenario
- The topology of this scenario
- Configuring the scenario
- Demonstration of the scenario

© Copyright IBM Corp. 2005. All rights reserved.

10.1 About this scenario

This scenario demonstrates the following:

- How to monitor the workload of a Web application by EWLM.
- How to set up and use Tivoli Provisioning Manager to start and stop a WebSphere Application Server.
- The value of such products in an heterogeneous environment.

The scenario starts with one WebSphere Application Server up and running. Then, using a workload generator, WebSphere Studio Workload Simulator, we overload the application Server. The monitoring is done by using EWLM. When EWLM reports the performance problem, we provision a new WebSphere Application Server in the Web application system. Then, we observe that the performance problem is gone by using the EWLM monitoring capabilities.

Figure 10-1 illustrates this scenario.

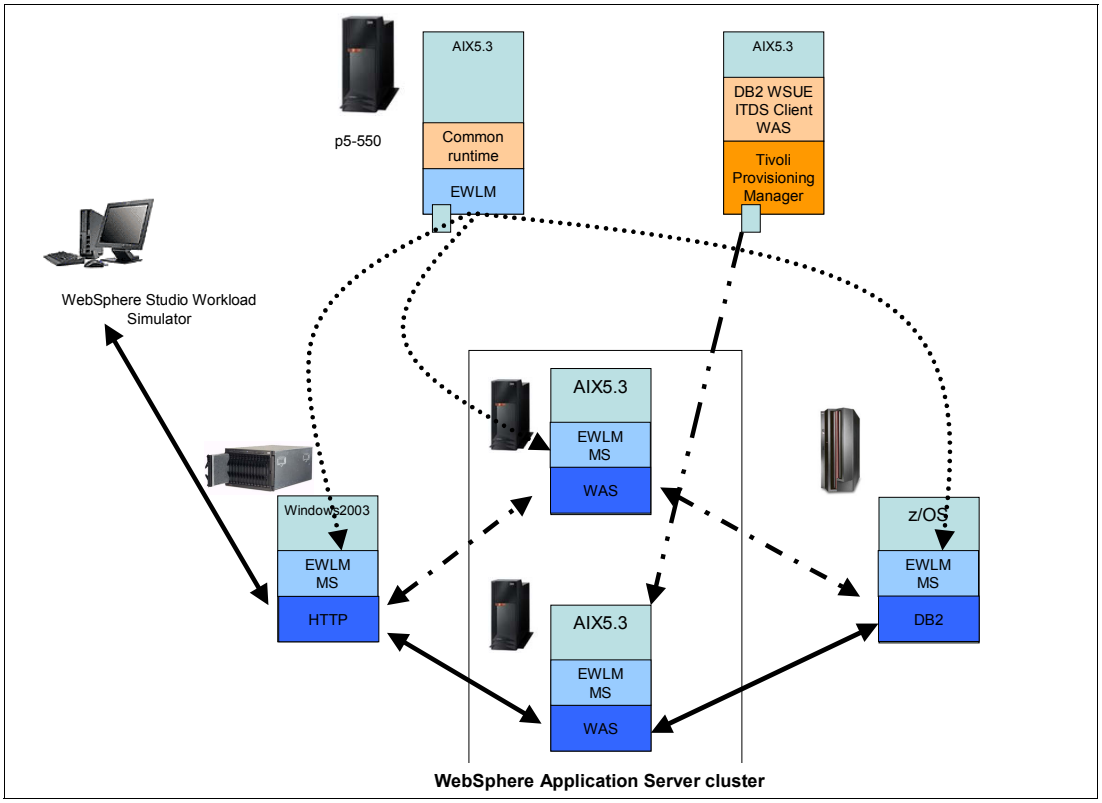

Figure 10-1 Scenario overview

Today, there is no integration between Tivoli Provisioning Manager and EWLM. We can monitor the performance of a Web application using EWLM and analyze the bottleneck of the Web application performance. Then, we can provision new resources using Tivoli Provisioning Manager, manually. For example:

- EWLM provides the information to find the bottle neck and to decide which resource we should provision.
- Tivoli Provisioning Manager simplifies the operation to set up new servers through workflows. When we build a workflow, only by executing the workflow, resources are provisioned without any platform dependent operations.

10.2 The topology of this scenario

In this scenario, we use a Web application called Trade3 which consists of an HTTP server, a WebSphere Application Server and a DB2 Universal Database server as shown in Figure 10-2. Trade3 is an application to demonstrate online trading service for stock broker.

Note: Trade3 is the third generation of the WebSphere end-to-end benchmark and performance sample application. For more information, refer to the following:

http://www-306.ibm.com/software/webservers/appserv/benchmark3.html

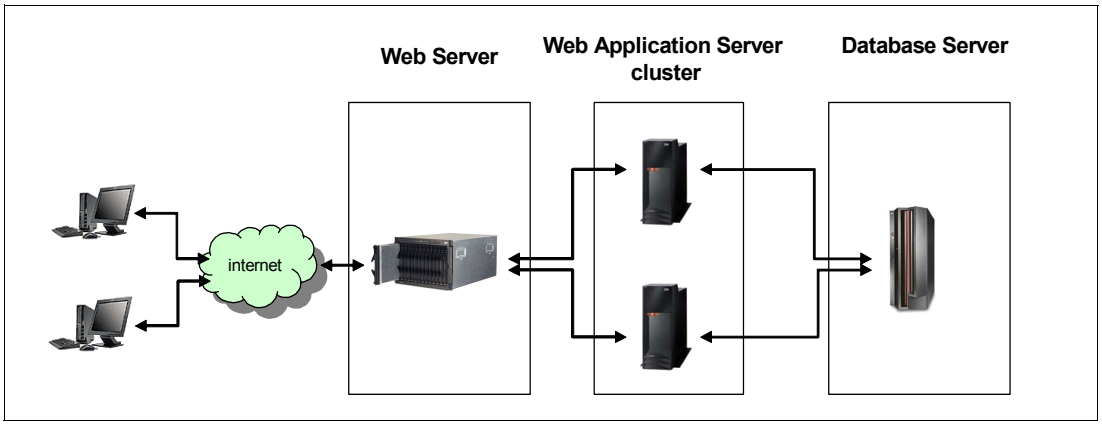

Figure 10-2 Trade3 application topology overview

The HTTP request received by the HTTP server is forwarded to the WebSphere Application Server. On the application server, the application accesses a database and the following occurs:

- The HTTP server runs on Windows 2003.
- Two WebSphere Application Server run on AIX 5.3.

  We configure the WebSphere Application Server nodes as a cluster and the WebSphere Application Server plug-in to point both servers. At the beginning of the scenario, only one WebSphere Application Server is started. The other server will be started by Tivoli Provisioning Manager when the response time becomes too high, beyond its performance goal.
- The DB2 server runs on z/OS.

Table 10-1 shows the network configuration and the role of each server. To stress the application, we use WebSphere Studio Workload Simulator.

Table 10-1 Scenario topology

Host Name	OS	Role
veaix41	AIX V5.3	EWLM Domain Manager
veaix42	AIX V5.3	Tivoli Provisioning Manager Server
veaix43	AIX V5.3	IBM Tivoli Directory Server
msnwin1	Windows 2003	Web Server - IBM HTTP Server - EWLM Managed Server
msnaix1	AIX V5.3	Web Application Server - WebSphere Application Server (Provisioned) - EWLM Managed Server
msnaix2	AIX V5.3	Web Application Server - WebSphere Application Server - EWLM Managed Server - WebSphere Application Server Deployment Manager
wtsc69	z/OS 1.6	Database server - DB2 Universal Database

In our scenario, we generate a high number of accesses to our Web application. We monitor it by EWLM and we provision a new WebSphere Application Server. In order to make our scenario simple, we do not install/uninstall WebSphere Application Server: we start the WebSphere Application Server and the EWLM Managed Server by using Tivoli Provisioning Manager.

Important: In this scenario, we use AIX nodes as management servers though we could use Windows or Linux servers as management servers as well. The management server functions of EWLM and Tivoli Provisioning Manager are independent from the operating systems.

10.3 Configuring the scenario

In this section, we describe how to configure our environment to run the scenario.

10.3.1 Setting up the EWLM environment

For the installation and the initial configuration of EWLM, refer to Part 2, “Installing the Virtualization Engine” on page 83.

We had to define a new domain policy which monitors our application’s transaction. For information about the configuration step of EWLM domain policy, refer to Chapter 5 of *IBM Enterprise Workload Manager*, SG24-6350.

10.3.2 Configuring Tivoli Provisioning Manager

In this section, we describe how to set up the Tivoli Provisioning Manager environment for our scenario. For more detailed information about to build a data center model, refer to *Provisioning On Demand Introducing IBM Tivoli Intelligent ThinkDynamic Orchestrator*, SG24-8888.

Adding resource pool and servers

During the installation process of Tivoli Provisioning Manager, a data center model sample is imported in the system. In our environment, we create a new data center model and we have to re-initialize the data center model by using the reinit.sh command located in the <TPM installed directory>/tools/ directory.

The **reinit.sh** command:

- Requires a datacentermodel.xml file as a parameter.
- Initializes everything in the data center model using the data center model which you specify.

Attention: The **reinit.sh** command deletes the current data center model and replaces it with the new data center model that you specify.

In the new data center model, we only define an Tivoli Provisioning Manager server and some basic devices such as fabric and switches. For a definition of the data center model, refer to 3.2.1, “The data center model” on page 41.

Example 10-1 describes the data center model that we use to create our configuration. We customize the datacentermodel.xml file that is provided with Tivoli Provisioning Manager. We change some value depending on the environment, such as the host name, and we add switch-fabric, switch, and software product. Then, we add two nic attributes called connected-to-switch and connected-to-port.

Example 10-1 Data center model

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- edited with XMLSPY v5 rel. 4 U (http://www.xmlspy.com) by IBM Corporation (IBM
Corporation) -->
<!DOCTYPE datacenter SYSTEM "xmlimport.dtd">
<datacenter>
   <switch-fabric name="ITSO_Fabric" />
   <switch name="ITSO_Switch" locale="en_US" fabric="ITSO_Fabric" failed="false"
in-maintenance="false">
      <switch-module>
         <switch-port failed="false" managed="false" netboot-enabled="false" vlan="1"
port-number="1" enabled="true" layer1-interface-type="unknown" />
         <switch-port failed="false" managed="false" netboot-enabled="false" vlan="1"
port-number="2" enabled="true" layer1-interface-type="unknown" />
         <switch-port failed="false" managed="false" netboot-enabled="false" vlan="1"
port-number="3" enabled="true" layer1-interface-type="unknown" />
         <switch-port failed="false" managed="false" netboot-enabled="false" vlan="1"
port-number="4" enabled="true" layer1-interface-type="unknown" />
         <switch-port failed="false" managed="false" netboot-enabled="false" vlan="1"
port-number="5" enabled="true" layer1-interface-type="unknown" />
         <switch-port failed="false" managed="false" netboot-enabled="false" vlan="1"
port-number="6" enabled="true" layer1-interface-type="unknown" />
         <switch-port failed="false" managed="false" netboot-enabled="false" vlan="1"
port-number="7" enabled="true" layer1-interface-type="unknown" />
```

```
         <switch-port failed="false" managed="false" netboot-enabled="false" vlan="1"
port-number="8" enabled="true" layer1-interface-type="unknown" />
      </switch-module>
      <switch-vlan vlan-number="1" />
   </switch>
   <subnetwork name="itso_ve" ipaddress="10.1.1.0" netmask="255.255.255.0">
      <vlan vlan-number="1" fabric="ITSO_Fabric" />
   </subnetwork>
   <software-category name="TIO" />
   <software name="TIO" locale="en_US" type="SOFTWARE" version="2.1" package_path="-"
category="TIO" />
   <spare-pool name="TPM Server" os-type="AIX" vlan="1" fabric="ITSO_Fabric">
      <server name="veaix42">
         <nic name="en0" macaddress="7241E0003002" netboot-enabled="false" managed="true"
connected-to-switch="ITSO_Switch" connected-to-port="3">
            <network-interface name="en0" ipaddress="10.1.1.187" netmask="255.255.255.0"
/>
         </nic>
         <property name="platform" value="9113-550"/>
         <sap name="scp client" port="22" host="false" app-protocol="SCP"
protocol-type="ipv4">
            <credentials search-key="primary" is-default="true">
               <password-credentials username="root"/>
            </credentials>
         </sap>
         <sap name="ssh service" port="22" host="true" app-protocol="SSH"
is-device-model="SSH Service Access Point" protocol-type="ipv4">
            <default-sap operation-type="execute-command"/>
            <credentials search-key="secondary" is-default="false">
               <password-credentials username="tioadmin" password="password"/>
            </credentials>
            <credentials search-key="primary" is-default="true">
               <rsa-credentials username="tioadmin"/>
            </credentials>
         </sap>
         <sap name="snmp-set" port="161" host="true" app-protocol="SNMP"
is-device-model="SNMP V1 Service Access Point" protocol-type="ipv4" context="write">
            <default-sap operation-type="set-attribute"/>
            <credentials search-key="primary" is-default="true">
               <snmp-credentials community="public"/>
            </credentials>
         </sap>
         <sap name="snmp-get" port="161" host="true" app-protocol="SNMP"
is-device-model="SNMP V1 Service Access Point" protocol-type="ipv4" context="read">
            <default-sap operation-type="get-attribute"/>
            <credentials search-key="primary" is-default="true">
               <snmp-credentials community="public"/>
            </credentials>
         </sap>
         <software-association software="AIX OS" state="running"/>
         <software-association software="TIO" state="running"/>
      </server>
   </spare-pool>
   <kanaha-config>
      <dcm-object id="0">
         <property component="DEPLOYMENT_ENGINE" name="PackageIntermediateHost"
value="veaix42"/>
         <property component="DEPLOYMENT_ENGINE" name="PackageRepositoryHost"
value="veaix42"/>
         <property component="DEPLOYMENT_ENGINE" name="PackageWebHost" value="veaix42"/>
```

```
        <property component="DEPLOYMENT_ENGINE" name="SoftwareRepositoryHost"
value="veaix42"/>
      </dcm-object>
   </kanaha-config>
</datacenter>
```

After the re-initialization, we add resource pools, as described in Table 10-2, and each server by using the Web GUI that is provided by Tivoli Provisioning Manager.

Table 10-2 Resource pools

Resource Pool	Description	Servers
Tivoli Provisioning Manager Server	Tivoli Provisioning Manager server resource pool	veaix42
Managed From AIX	A resource pool for Managed From AIX Tivoli Provisioning Manager server which running WebSphere Application Server	msnaix2
SparePool	A resource pool for servers which are ready to start WebSphrer Application Server.	msnaix1

To add resource pool, click **Resource Pool** on the left navigation panel, and click **Edit**.

- To add servers in each resource pool:
  a. Click the resource pool name in the servers menu on the left navigation panel to show servers which belongs to the resource pool.
  b. Click **Edit** on the resource pool window and add new server.
  c. Then, click the server name. In the General tab in the server window, add NIC and network interface. When you create NIC, do not forget to fill in the switch, switch port, and VLAN to which the NIC belongs.

Creating credentials on each server

To execute WebSphere Application Server startup workflows, we need to define Service Access Points to transfer files and to execute commands. We create SCP access points for transferring files and SSH access points for executing commands.

We generate an RSA keyfile for these Service Access Points. Figure 10-3 on page 250 shows the Services Access Points that we have defined. Figure 10-4 on page 250 shows the workflows that we defined for each Service Access Point.

We select SSH Service Access Point device driver for each Service Access Point. In order to display this window, select the link for Service Access Points. Then, set SCP as the default access point of file-transfer device operation and SSH for execution-command device operation.

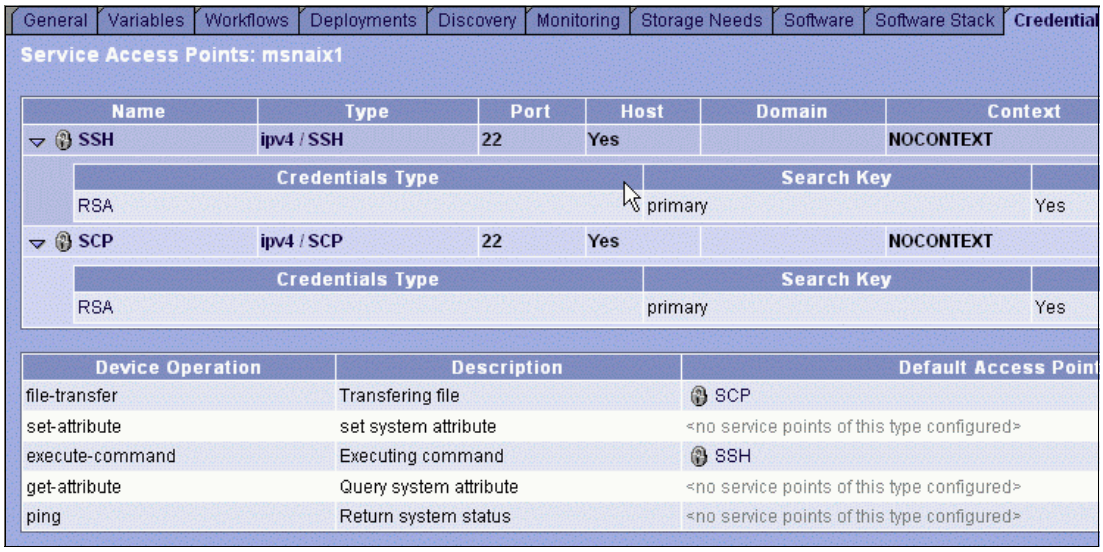

Figure 10-3 Service Access Points

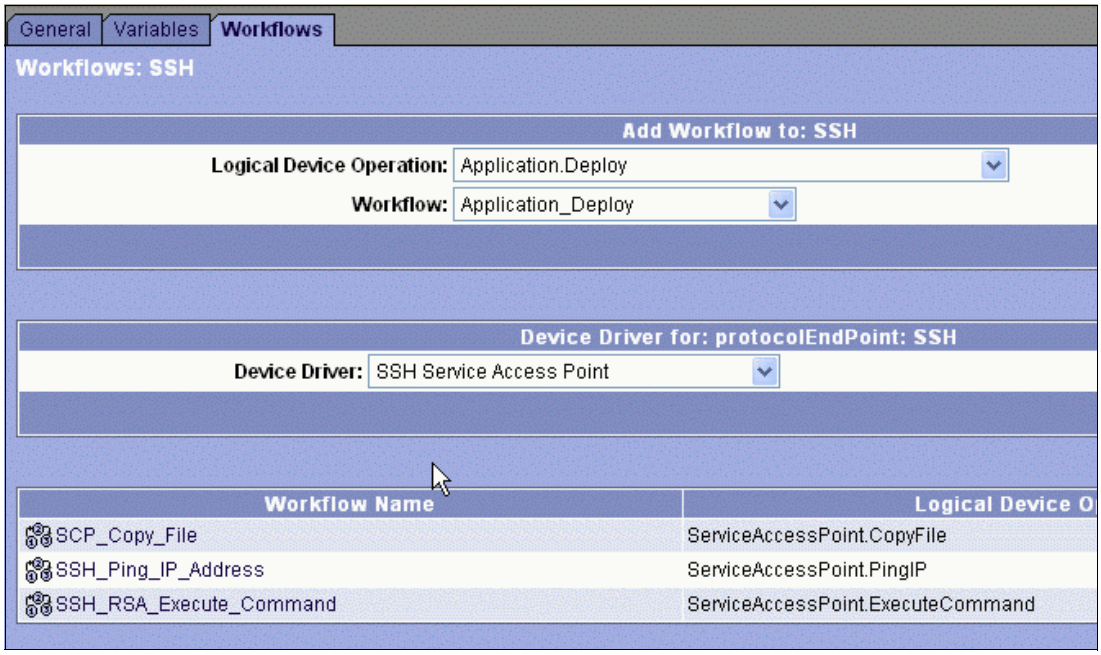

Figure 10-4 The workflow definition of Service Access Point

For information about establishing SSH connection between servers, refer to *Tivoli Provisioning Manager Installation Guide for UNIX*, GC32-1615.

Assigning initialization workflows for each resource pool

The initialization workflow moves the server to the target resource pool and initializes them. If we configure a software stack on the target resource pool, the initialization workflow tries to install software products as defined in the software stack.

In our scenario, we do not install WebSphere Application Server and EWLM Managed Server when we provision a server (before the products are already installed). Therefore, we use this workflow only to move the server to the target resource pool.

To add this workflow in the Resource Pool:

1. Go to the Workflow tab on the ResourcePool view.
2. Select **SparePool.InitializeServer** as the Logical Device Operation.
3. Select **Default_Pool_Initialize_Server** as the Workflow.
4. Click **Add Selected Workflow**.

Customizing and creating workflows

The WebSphere Application Server workflow provided with Tivoli Provisioning Manager V2.1 only supports Windows platform. Therefore, in order to provision a WebSphere Application Server on AIX, we need to customize the workflow.

In our scenario, we use a start and stop workflow for the WebSphere Application Server. We changed the scriptlet in the start or stop workflow that starts and stops the application server on AIX and saved them as IBM_WAS_Start_Custom and IBM_WAS_Stop_Custom workflows.

The EWLM Managed Serve workflow is not provided with Tivoli Provisioning Manager and, at the time of writing, there was not a workflow available on the IBM Ondemand Automation Catalog site (OPAL). Therefore, we need to create a brand new workflow to start or stop the EWLM Managed Server.

The used workflows implement the Software.start and Software.stop logical device operation. A logical device operation provides a common interface to the managed data center model which is independent from the type of devices and software. For example, the implementation of the WebSphere Application Server start workflow and the EWLM start workflow is different, but both workflows are called by the deployment engine via the same interface called Software.Start.

In addition, we have to create dummy_install and dummy_uninstall workflows which implement Sofware.Install and Software.Uninstall logical operations. These workflows do not execute anything. However, they are called from the resource pool initialization workflow and need to be there.

Note: Although starting EWLM Managed Server in a background process is not supported, we used this configuration in our scenario for functional purpose, not with the intent to put it in production.

Creating software product

The Software Product WebSphere Application Server is predefined as WebSphere Application Server 5.1. However, we create new software product for our environment. We change properties, variables, and workflows of the Software Product as follows:

- Variables: We need to specify the was.service.name variable, which is used for the starting server workflow, and the was.responsefile.template.name, which is used for the installing software workflow.

  In our scenario, WebSphere Application Server is already installed. Therefore, we set only was.service.name.

- Workflows: We add the dummy_install, dummy_unistall, IBM_WAS_Start_Custom, and IBM_WAS_Stop_Custom workflows as created previously.

In addition, we create a new software product called EWLM Managed Server. We define a ewlm.dmport variable which indicates the EWLM Domain Manager port, and we add the dummy install and uninstall and start and stop EWLM workflows.

Defining a software stack and associate it with a resource pool

To execute the initialization workflow, we need to configure a software stack and to associate it with the target resource pool as follows:

- The software stack we created consists of the EWLM Managed Server and the WebSphere Application Server software product we have created.
- To associate it with the target resource pool, go to the Software Stack tab of the target resource pool and select the software stack configuration. Then, click **Save**.

During the initialization process, the installation workflow of each software product is executed. In our environment, the installation process does not happen because we define dummy_install workflow which does not execute anything.

Setting servers

The workflow to start the WebSphere Application Server requires to set the was.installLocation variable which points to the directory on which the WebSphere Application Server is installed (target server: msnaix1). To set this variable, go to the Variables tab on the server, enter the name and the value, and then click **Add**.

In addition, the ewlm.workdir and ewlm.installLocation variables are used by our custom EWLM startup workflow. Figure 10-5 shows our variable settings for the servers.

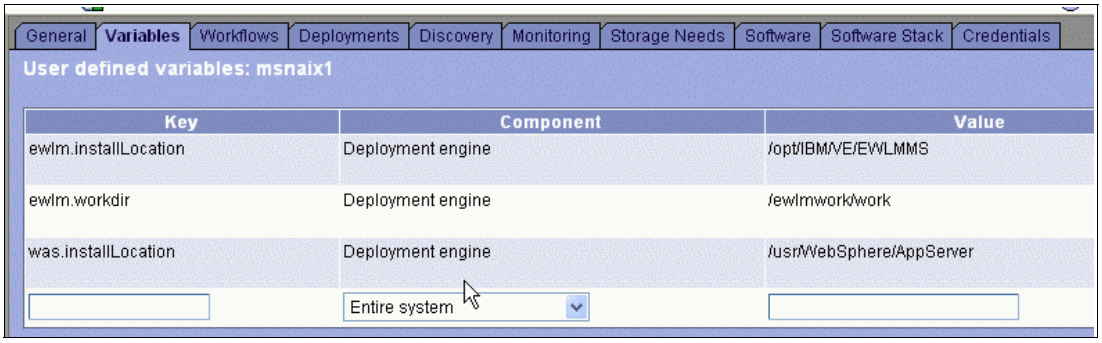

Figure 10-5 Setting server variables

Lastly, we associate the software product with the server. Go to the Software tab and choose the software product. Then, click **Associate**.

10.4 Demonstration of the scenario

The steps of this scenario are as follows:

1. When we start, we have one HTTP server, one WebSphere Application Server and one DB2 Universal Database server. Another WebSphere Application Server which is named msnaix1 is configured and booted. However, this WebSphere Application Server is not activated.

   The EWLM Managed Server is running on each active server (not on the msnaix1 server). To confirm that the EWLM works correctly, log in to the Control Center of EWLM and show the managed server status by clicking **Managed Server** in the list of Management on the right navigation panel.

As shown in Figure 10-6, the status of msnaix1 is `Communication error` because the EWLM Managed Server is down and the status of the other servers is Active. Active service policy is the service policy which is activated on the server. We use ITSO Trade3 and Plants Service Policy as the service policy.

Select	Name	State	Active service policy	Operating system
○	msnaix1.itso.ibm.com	Communication error		AIX
○	msnaix2.itso.ibm.com	Active	ITSO Trade3 and Plants Service Policy	AIX
○	msnwin1	Active	ITSO Trade3 and Plants Service Policy	Windows
○	wtsc69.itso.ibm.com	Active	ITSO Trade3 and Plants Service Policy	z/OS

Page 1 of 1 Total: 4 Filtered: 4 Displayed: 4

Figure 10-6 Managed Server status before provisioning a server

In our environment, we have different service classes which are for portfolio calculation, login, trade action and so on. We defined the goal for each service class.

The achievement rate is indicated by the Performance Index (PI). PI indicates how well a service class is attaining the goal defined for it. If PI is grater than 1.00, it means that the goal is not achieved.

2. When we start WebSphere Studio Workload Simulator, it generates a lots of user sessions. After a while, we got PI as shown in Figure 10-7.

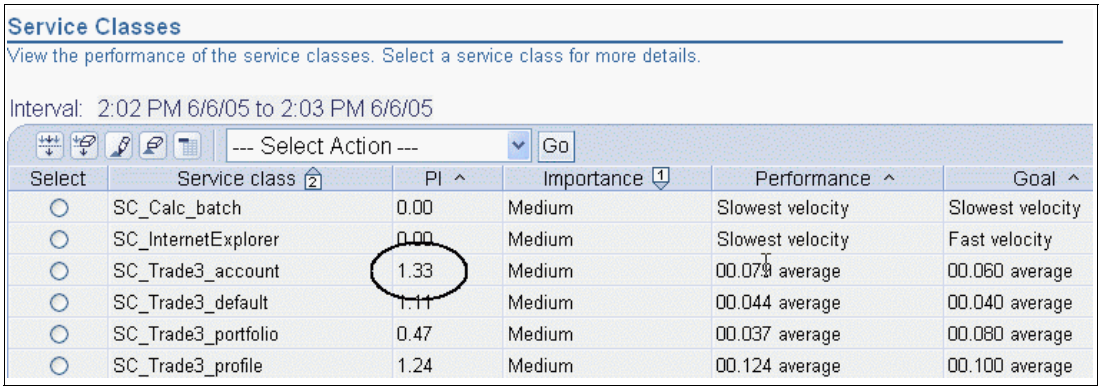

Service Classes

View the performance of the service classes. Select a service class for more details.

Interval: 2:02 PM 6/6/05 to 2:03 PM 6/6/05

--- Select Action --- Go

Select	Service class	PI	Importance	Performance	Goal
○	SC_Calc_batch	0.00	Medium	Slowest velocity	Slowest velocity
○	SC_InternetExplorer	0.00	Medium	Slowest velocity	Fast velocity
○	SC_Trade3_account	1.33	Medium	00.079 average	00.060 average
○	SC_Trade3_default	1.11	Medium	00.044 average	00.040 average
○	SC_Trade3_portfolio	0.47	Medium	00.037 average	00.080 average
○	SC_Trade3_profile	1.24	Medium	00.124 average	00.100 average

Figure 10-7 Service Class performance

As shown in Figure 10-7, we can find that some service classes cannot satisfy their performance goal.

3. In order to figure out the bottleneck, we use the EWLM application topology of one of those service classes. We select **Application Topology** in the Select Action drop-down menu, and click **Go**. If we open the Application Topology, we get the topology information as shown in Figure 10-8.

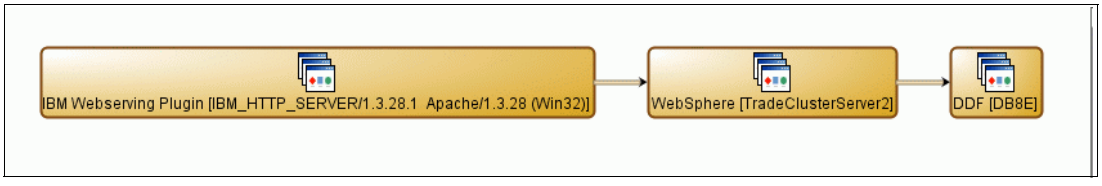

Figure 10-8 Application topology of SC_trade3_account service class

As shown in Figure 10-8, we have three servers running for this application, as expected. Arrows between servers indicate the flow of the transaction of the application.

4. Click **Table View** on the menu bar of the Application Topology view to get more detailed information for the application. As shown in Figure 10-9, there is some process delay for the WebSphere Application Server msnaix2. Therefore, we expect that the processing on msnaix2 causes the mis-achievement of the PI.

Hop	Name	Type	Platform	Average response time ...	Average active time (se...	Processor time (secon...	Processor delay (%
0	IBM Webserving Plugin [IBM_HTTP...	Application		00.080	00.004	00.0	0%
0	msnwin1	Server	Windows	00.080	00.004	00.0	0%
0	MSNWIN1/PID=0000002600	Instance		00.080	00.004	00.0	0%
1	WebSphere [TradeClusterServer2]	Application		00.070	00.044	01.950355	6%
1	msnaix2.itso.ibm.com	Server	AIX	00.070	00.044	01.950355	6%
1	msnaix2.TradeClusterServer2	Instance		00.070	00.044	01.950355	6%

Figure 10-9 Detailed informations of transaction

5. We can obtain the processor utilization information for the servers from EWLM.

   Click **Managed Server** in the Monitor menu list on the left navigation panel. The Managed Server status view comes up as depicted in Figure 10-10. At this point, we prove that the processor utilization of server msnaix2 is very high.

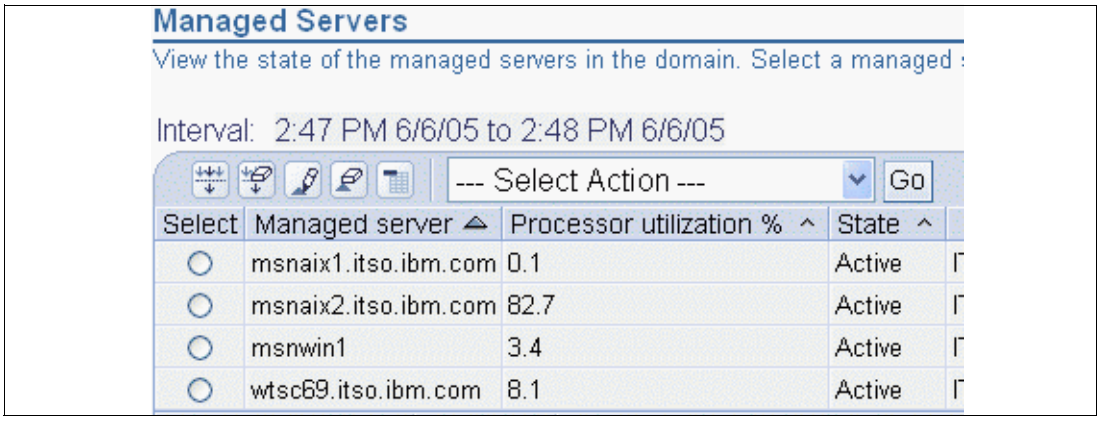

Figure 10-10 CPU Utilization of each node

6. We start a new WebSphere Application Server using Tivoli Provisioning Manager.

   Figure 10-11 shows our data center model graphically. From the left of figure, Managed From AIX, SparePool, and Tivoli Provisioning Manager Server resource pools are displayed. Each small blue box depicts one server.

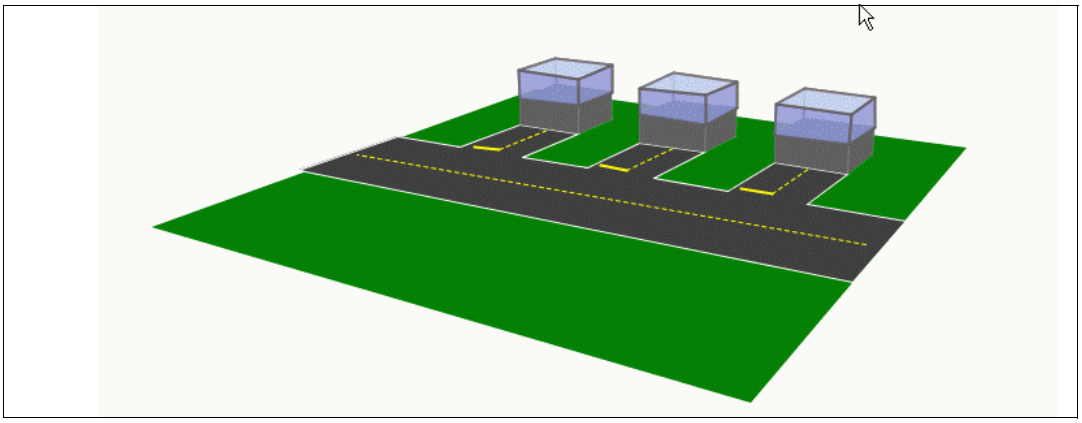

Figure 10-11 Data center model before provisioning

7. We move msnaix1 from the SparePool to the Managed From AIX resource pool. If we select the Managed From AIX pool, as shown in Figure 10-12, the Default_Pool_initialize_Server workflow, which is defined in Managed From AIX resource pool, is called automatically.

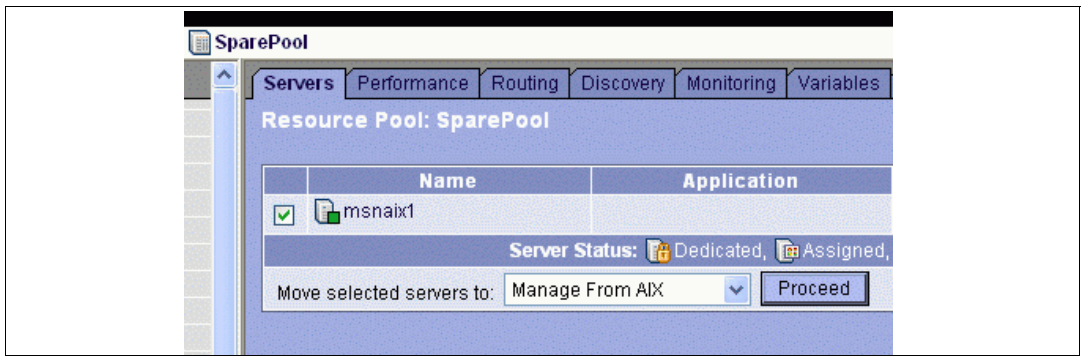

Figure 10-12 Moving server to Managed From AIX resource pool

If we define the installation workflow for the Software Product, which is in Software Stack, the installation process is initiated by the initializing workflow. In our environment, we use the dummy_install workflow, which does nothing. Therefore, the initialization process moves the server but does not install the software.

Figure 10-13 shows the data center model after executing the initialization workflow.

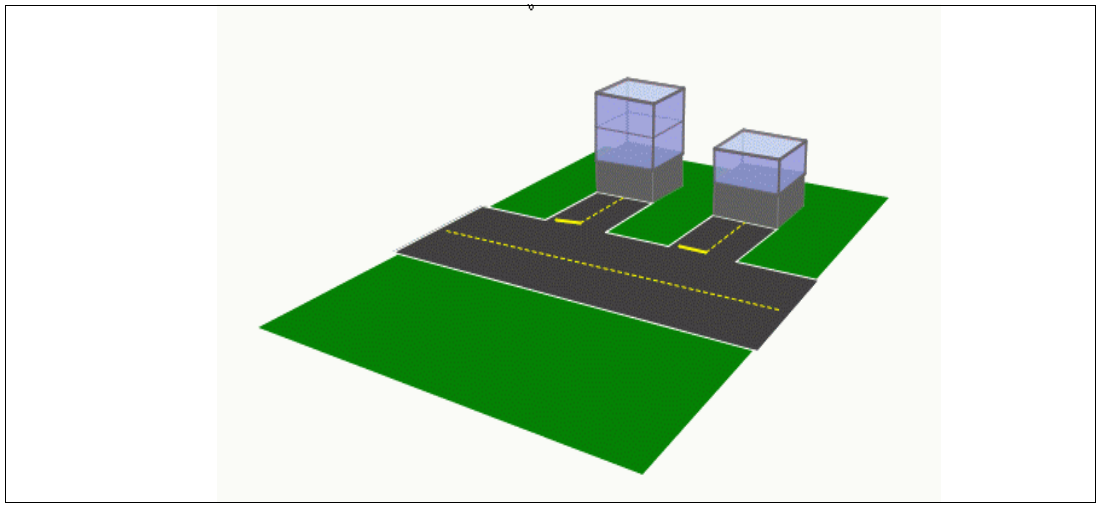

Figure 10-13 Data center model after provisioning

We successfully added msnaix1 in the Managed From AIX resource pool.

8. We start the EWLM Managed Server and the WebSphere Application Server via the defined workflows. The software installation workflow requires the DeviceID of the server and the SoftwareID of the software. Figure 10-14 shows the execution of this workflow.

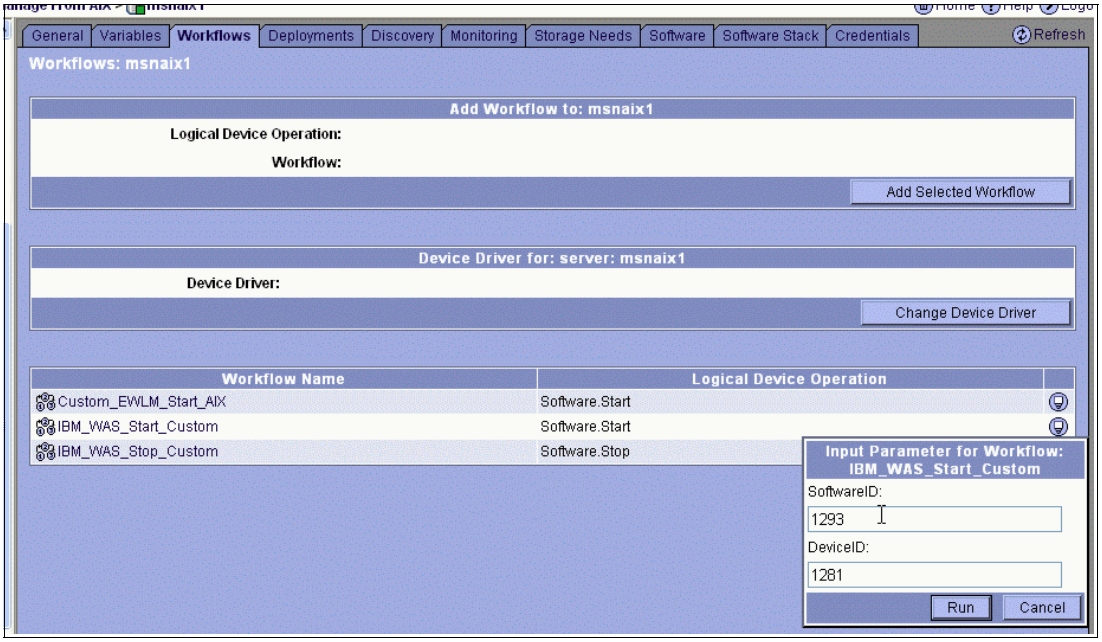

Figure 10-14 Running the workflow to start EWLM and WebSphere Application Server

When we run a workflow, we redirect to the status window of the initiated workflow. You can check the status of workflow. The status are the following:

`created`	The workflow is just initiated.
`in-progress`	The workflow is running.
`failed`	The workflow is failed
`success`	The workflow is succeeded.

Figure 10-15 shows the status window of the workflow.

Figure 10-15 Status of Workflow

If you want to see more detailed workflow process or to confirm the status later, you can check it on the Deployment Requests window. Figure 10-16 shows the location of the Deployment Requests window.

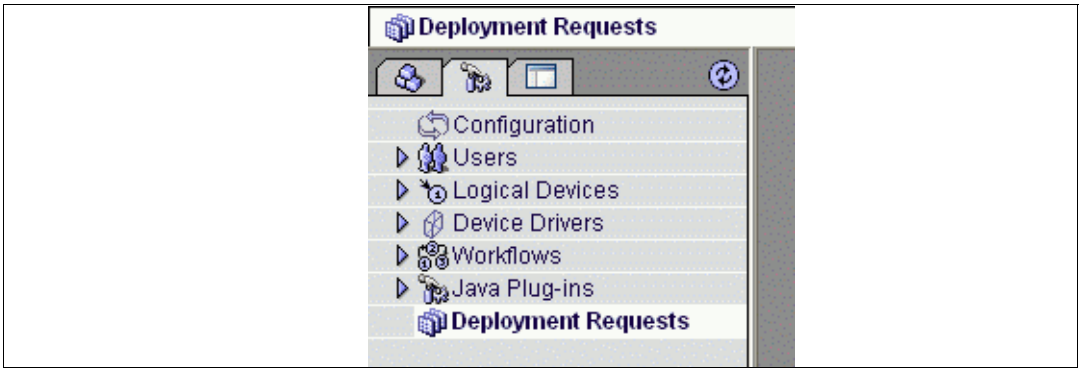

Figure 10-16 Deployment Requests window

9. When we succeed to execute the workflow, the EWLM Managed Server and the WebSphere Application Server are started on msnaix1. If the EWLM Managed Server is correctly started, the EWLM Domain Manager and the EWLM Managed Server start the communication automatically so that it might appear in Figure 10-17. The status of msnaix1 is changed automatically to Active.

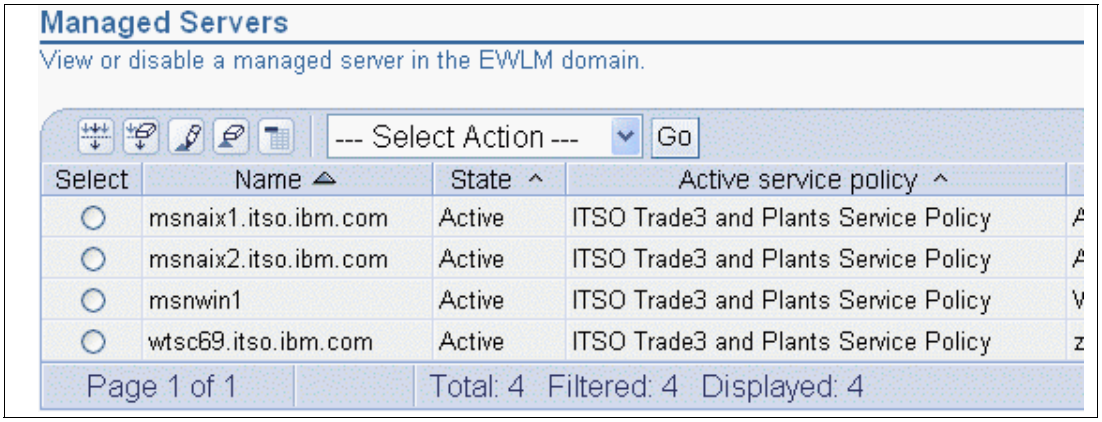

Figure 10-17 Managed Server status on EWLM

10. Then, the EWLM Managed Server on the server msnaix1 reports about its transaction performance. At this point, we check Service Class performance again. Figure 10-18 shows the status of the Service Class after provisioning the server msnaix1.

--- Select Action --- Go

Select	Service class	PI	Importance	Performance
○	SC_Calc_batch	0.00	Medium	Slowest velocity
○	SC_InternetExplorer	0.00	Medium	Slowest velocity
○	SC_Trade3_account	0.85	Medium	00.050 average
○	SC_Trade3_default	0.74	Medium	00.029 average
○	SC_Trade3_portfolio	0.35	Medium	00.028 average
○	SC_Trade3_profile	0.95	Medium	00.094 average
○	SC_Trade3_quotes	0.57	Medium	00.017 average
○	SC_Trade3_shopping	0.79	Medium	97% in 00.300
○	SC_xPlants_default	0.00	Medium	00.000 average
○	SC_xPlants_login	0.00	Medium	0% in 00.150
○	SC_xPlants_shopping	0.00	Medium	0% in 00.200
○	EWLM Service Class	<None>	<None>	<None>
○	SC_Trade3_general	<None>	<None>	<None>
○	SC_xPlants_general	<None>	<None>	<None>

Page 1 of 1 Total: 14 Filtered: 14 Displayed: 14

Figure 10-18 Service class performance

11. Now, all the PIs are under 1.00 and the performance requirements of each service class are satisfied. To confirm the transaction flow, we go to the Application Topology view again. Figure 10-19 shows the new application topology after the provisioning.

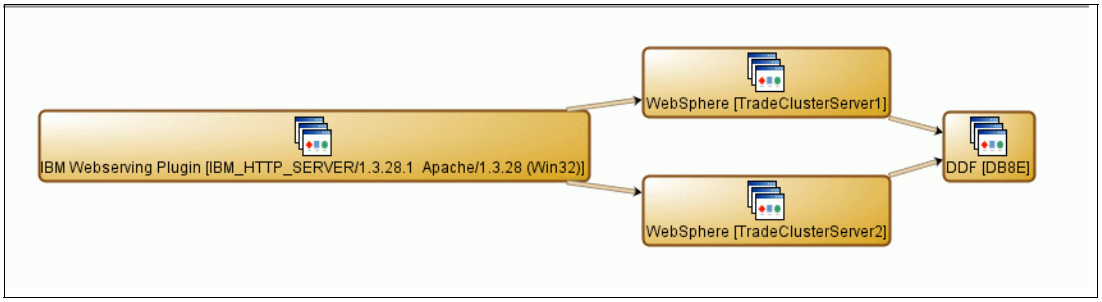

Figure 10-19 Application topology after provisioning server

The application topology is updated automatically by EWLM. EWLM gets the transaction information from the EWLM Managed Server and recreates the application topology automatically.

At this point, we check the processor utilization for all the servers again. Figure 10-20 shows the status of these servers. We can find that the processor utilization of the server msnaix2 is going down and msnaix1 is going up.

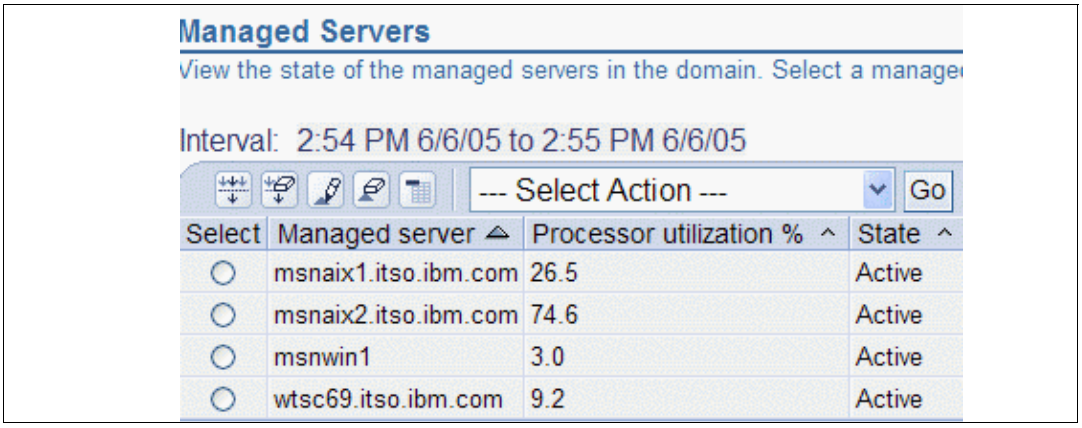

Figure 10-20 CPU utilization after provisioning server

Thus, we can confirm that the provisioning of WebSphere Application Server is successful.

11

How to monitor an application in a heterogeneous environment using the VE console and EWLM

This chapter describes how we used the Virtualization Engine components to monitor our heterogeneous environment. It includes the following topics:

- The topology for this scenario
- Configuring the heterogeneous scenario
- Validating the configuration
- Testing the heterogeneous scenario
- Conclusions

© Copyright IBM Corp. 2005. All rights reserved.

11.1 The topology for this scenario

In Part 2, “Installing the Virtualization Engine” on page 83, we describe the installation process for the Virtualization Engine components in a platform specific environment. Specifically, we had three topologies that we tested:

- An all-Windows topology
- An all-Linux topology
- An all-AIX environment (with the exception of the IBM Director which ran on another platform)

Within each of the topologies, both the management servers and managed servers ran under the same operating system. We detailed this information as each platform had certain idiosyncrasies. However, when the Virtualization Engine components get installed and configured, the use and implementation is the same, regardless of the chosen platform.

In this chapter, we highlight how we used the Virtualization Engine components to monitor our workload running across a heterogeneous environment, regardless of the operating system on which the Virtualization Engine components or the workload is running.

We used the EWLM Domain Manager, the Virtualization Engine Console, the IBM Director, and the Tivoli Provisioning Manager to monitor and manage our environment. We also had multiple managed servers running on multiple operating systems hosting one HTTP server, two WebSphere Application Server servers, and one database server. Let’s discuss our environment in more detail.

11.1.1 Management servers

Our management servers consisted of a combination of AIX and Windows systems. As shown in Figure 11-1, we used a p5-550 with three LPARs running AIX 5.3 and an IBM BladeCenter HS20 running a Windows 2003 server as our management servers.

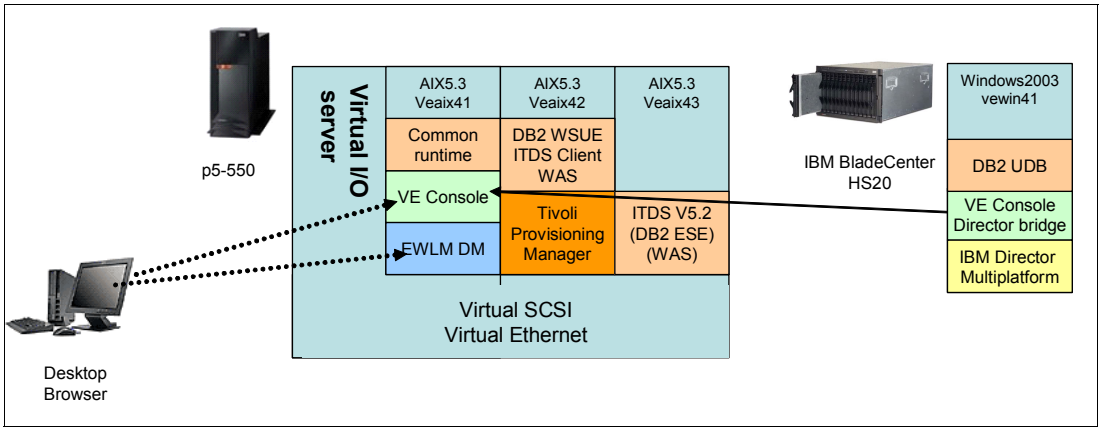

Figure 11-1 Management server configuration

We defined the management server configuration as follows:

- P5-550 LPAR 1: VEAIX41
  - Virtualization Engine Console 1.1
  - EWLM Domain Manager 1.1
  - Common runtime components: DB2 8.1, WebSphere Application Server 5.1, IBM Tivoli Directory Server 5.1
- P5-550 LPAR 2: VEAIX42
  - Tivoli Provisioning Manager 2.1
  - Tivoli Provisioning Manager prerequisites: WebSphere Application Server 5.1, DB2 WSUE 8.1, and IBM Tivoli Directory Server client 5.2
- P5-550 LPAR3: VEAIX43
  - IBM Tivoli Directory Server 5.2 (or LDAP) in support of Tivoli Provisioning Manager
  - IBM Tivoli Directory Server prerequisites: DB2 ESE and WebSphere Application Server
- HS20: VEWIN41
  - Virtualization Engine bridge 1.1 to IBM Director Multiplatform
  - IBM Director 4.2
  - IBM Director prerequisites: DB2 UDB 8.1

11.1.2 Managed servers

We set up a heterogeneous managed server environment to run the Trade3 workload consisting of Windows, AIX and z/OS.

- We configured all of the servers to be part of the EWLM management domain.
- We configured the Windows and AIX servers to be IBM Director managed nodes.

Restriction: The IBM Director Agent is not supported today on a z/OS, therefore it cannot be an IBM Director managed server.

- We configured one HTTP Server to run on the Windows node, two WebSphere Application Servers to run on separate AIX nodes, and a DB2 server to run on a z/OS node. The managed server topology as well as the middleware layout is shown in Figure 11-2.

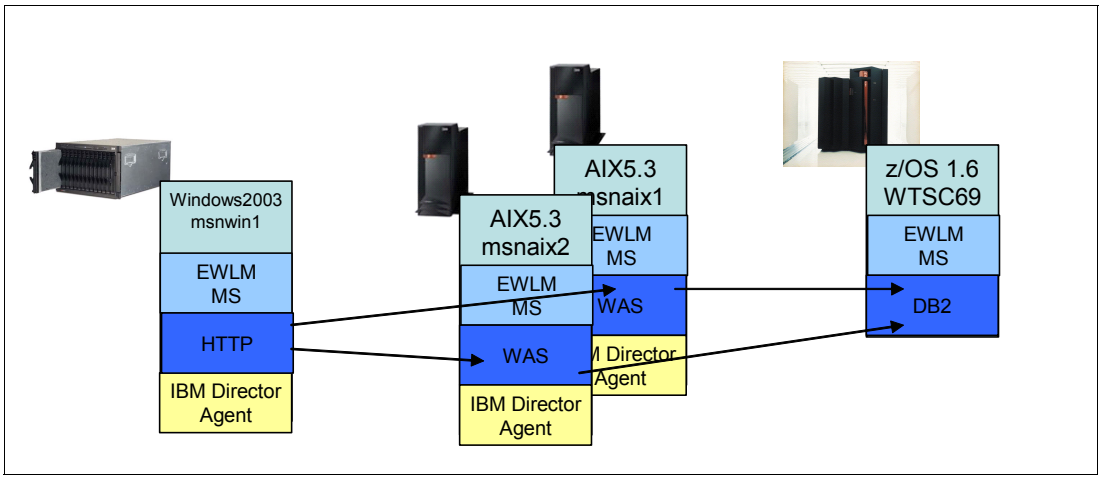

Figure 11-2 Managed server environment

We defined the managed server configuration as follows:

- HS20: MSNWIN1 (windows)
  - EWLM 1.1managed server
  - IBM Director Agent 4.2
  - IBM HTTP 1.3.28 server
- P5-550 LPAR 4 msnaix2:
  - EWLM 1.1 managed server
  - IBM Director Agent 4.2
  - WebSphere Application Server 5.1.1
  - WebSphere Application Server for Network Deployment 5.1.1
- P5-550 LPAR 5 msnaix1:
  - EWLM 1.1 managed server
  - IBM Director Agent 4.2
  - WebSphere Application Server 5.1.1
  - IBM HTTP 1.3.28 server 1.3.28
- z990 LPAR with z/OS 1.6: WTSC69
  - EWLM 1.1 managed server
  - DB2 8.2 server

11.1.3 The workload

We ran a workload simulator (WebSphere Studio Workload Simulator) on another system to drive the Trade3 transactions into the HTTP server to simulate our customer environment. We verified that the workload simulator was working correctly by verifying that the simulator was able to send the transactions to the HTTP server and not receive an incorrect response

The Trade3 application can be accessed from a Browser to buy or sell stocks, maintain user accounts, and user portfolios. The Trade3 workloads consist of an end-to-end Web application and a full set of Web primitives. It can be used as an end-to-end benchmark and performance sample application.

Figure 11-3 shows the entire environment for testing the Trade3 workload, which includes the management servers, the managed servers, and the workload simulator.

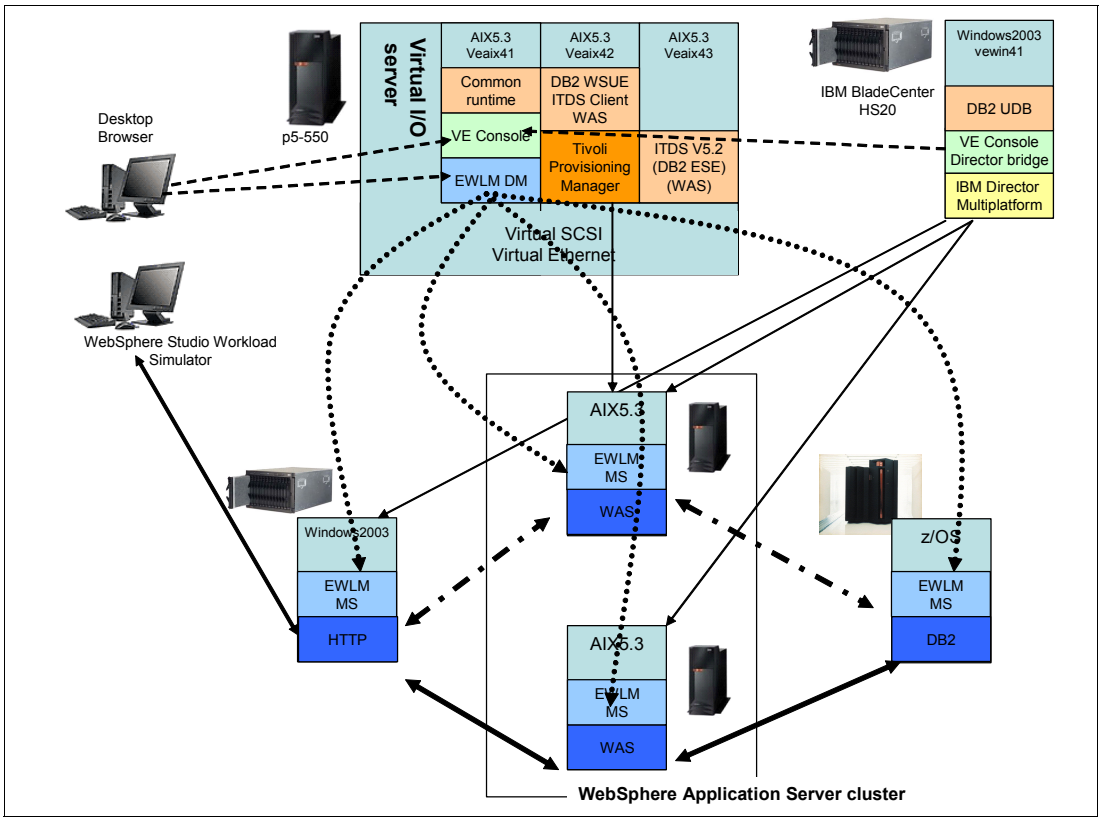

Figure 11-3 Entire heterogeneous environment

We used a browser to access the Virtualization Engine console and the EWLM control center. We used VNCviewer and SSH to access the Windows and AIX nodes and IBM personal communications to access the z/OS node.

11.2 Configuring the heterogeneous scenario

Our goal in this scenario was to use the EWLM control center and the Virtualization Engine Console to monitor the environment and Tivoli Provisioning Manager to provision a new server, when our goals were not being met. We used the environment described Figure 11-3 to run the Trade3 workload and to monitor it.

We divided the scenario into three parts:

- The first part of the scenario uses the EWLM control center and the Virtualization Engine Console to monitor the environment running under *normal* conditions. It starts with the EWLM performance goal being met and all of the systems are running without problems or exceptions. In this part, the workload is running with one HTTP Server, one WebSphere Application Server and one database manager, DB2. No other work was running on any of the systems except for the z/OS node.
- The second part of the scenario created a constraint on the WebSphere Application Server system which caused the performance goals not being met. We used the EWLM control center and the Virtualization Engine Console to monitor the systems and to

determine that the performance goals were not being met and to detect where the bottleneck was.

- In the third part of the scenario, we used Tivoli Provisioning Manager to provision a second WebSphere Application Server into the TradeCluster on an AIX node which was also provisioned into the EWLM management domain. We used the EWLM control center and the Virtualization Engine Console to monitor the effects of the provisioned system to verify that the environment is again running under *normal* conditions.

We walk through the functions of each of the products and how we used them to monitor our environment to verify that our environment had been dynamically updated.

11.2.1 EWLM consideration

An important aspect of EWLM is that all data collection and aggregation activities are driven by a common service level policy, called the EWLM *Domain Policy*. This policy is built by an administrator to describe the various business processes and the performance objectives for each process. See 3.1.2, “The domain policy” on page 37 for detailed information about setting up a Domain Policy.

Note: We used the domain policy which was created for *IBM Enterprise Workload Manager,* GC24-6350, with some minor modifications for the z/OS definition and alterations of the response time goals.

We set up an EWLM domain policy so that we could classify, prioritize, and monitor our Trade3 workload. Specifically, we did the following:

- Defined the platforms on which we were running.
- Qualified the Trade3 transactions into service classes.
- Associated a performance goal to each of the service classes.

You can see a copy of our domain policy in Appendix C, “Sample EWLM domain policy” on page 377.

The performance goal defined in EWLM is what we used to monitor our workload to verify that the defined environment was adequate for achieving our goals. We will discuss the process of verifying these goals as we go through our heterogeneous scenario in 11.4.1, “Monitoring the workload running under normal conditions” on page 281.

The EWLM configuration consisted of one Domain Manager and four managed servers on multiple operating systems. The EWLM domain policy was created with the definitions for Trade3 so that we could monitor our performance goals using the EWLM control center.

Before starting the scenario, we verified the following:

- EWLM was working correctly by verifying that all the managed servers appeared in the Managed Servers view and the transactions were increasing when the simulator was running.
- The HTTP server distributed the transactions to both WebSphere Application Servers and the WebSphere Application Servers sent the transaction request to DB2 by looking at the Application topology view.

11.2.2 IBM Director consideration

IBM Director is a very powerful tool for systems management. There are many functions which the IBM Director can perform such as event logging, hardware status monitoring, network configuration and resource threshold monitoring. To see how many of the functions that can be defined, refer to 3.3, “IBM Director Multiplatform” on page 44. A few of the resource monitors which can be set are identified in Figure 11-4.

We focused on a couple of the tasks: hardware status and resource monitors.

Note: The list of available resource monitors will vary between operating systems. For example, if the operating system is Windows, in addition to the resource monitors listed in Example 11-4, you are able to monitor CIM, Windows Device, performance, and service monitors.

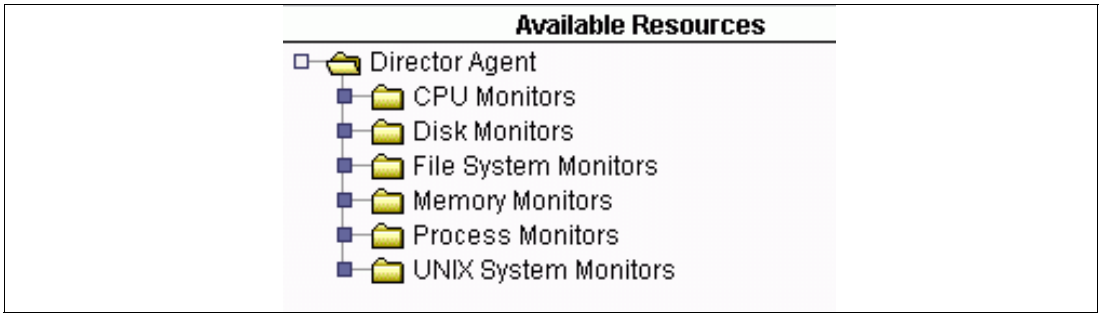

Figure 11-4 IBM Director Multiplatform resources to monitor

In our environment, we decided that the critical areas to monitor using the IBM Director were:

- The processor of the system.
- The percentage of disk space free
- The memory usage for the two WebSphere Application Server systems and the HTTP server

We set up thresholds so that we would be notified if those thresholds were being met as they might affect our workload and perhaps our performance goals. Some of our available thresholds are shown in Figure 11-5 which we subsequently added to the Virtualization Engine Console as described in 11.2.3, “Virtualization Engine console consideration” on page 268. Note that these thresholds are only a small number of the potential thresholds that can be monitored. Evaluate your own requirements, and set the respective thresholds accordingly.

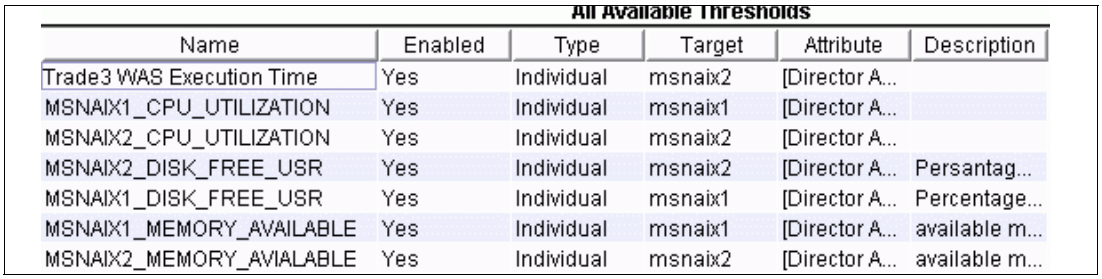

Name	Enabled	Type	Target	Attribute	Description
Trade3 WAS Execution Time	Yes	Individual	msnaix2	[Director A...	
MSNAIX1_CPU_UTILIZATION	Yes	Individual	msnaix1	[Director A...	
MSNAIX2_CPU_UTILIZATION	Yes	Individual	msnaix2	[Director A...	
MSNAIX2_DISK_FREE_USR	Yes	Individual	msnaix2	[Director A...	Persantag...
MSNAIX1_DISK_FREE_USR	Yes	Individual	msnaix1	[Director A...	Percentage...
MSNAIX1_MEMORY_AVAILABLE	Yes	Individual	msnaix1	[Director A...	available m...
MSNAIX2_MEMORY_AVIALABLE	Yes	Individual	msnaix2	[Director A...	available m...

Figure 11-5 IBM Director available thresholds

Note: Threshold amounts vary from operating system to operating system and from customer to customer. So, we did not describe the values that we set as it would probably not be relevant for most customers. Some systems, such as z/OS, operate efficiently at 100% processor use, while others do not. Therefore, the values of a threshold can be vastly different across platforms and customers. The middleware that is running on one of the platforms might also be more dependent on the processor than another.

Summary

The IBM Director configuration was set up to monitor processor, disk space, and available memory. Thresholds were set which would help to indicate potential warning and error conditions. Although there were many other definitions which we could have defined, our goal was to use the EWLM control center and Virtualization Engine Console to monitor the environment, which these definitions would satisfy.

11.2.3 Virtualization Engine console consideration

The Virtualization Engine console needs to configure the following:

- User IDs and groups
- Management sources

Customizing preferences and managing users and groups

We can customize some part of Virtualization Engine console such as rows in a table, refresh rate and so on. For example, Figure 11-6 shows the general parameter which we can customize. We can display these parameters by clicking **General** in Preferences on the left navigation pane. *Dashboad* is for preferences of dashboad, and *Charting* is preferences of chart.

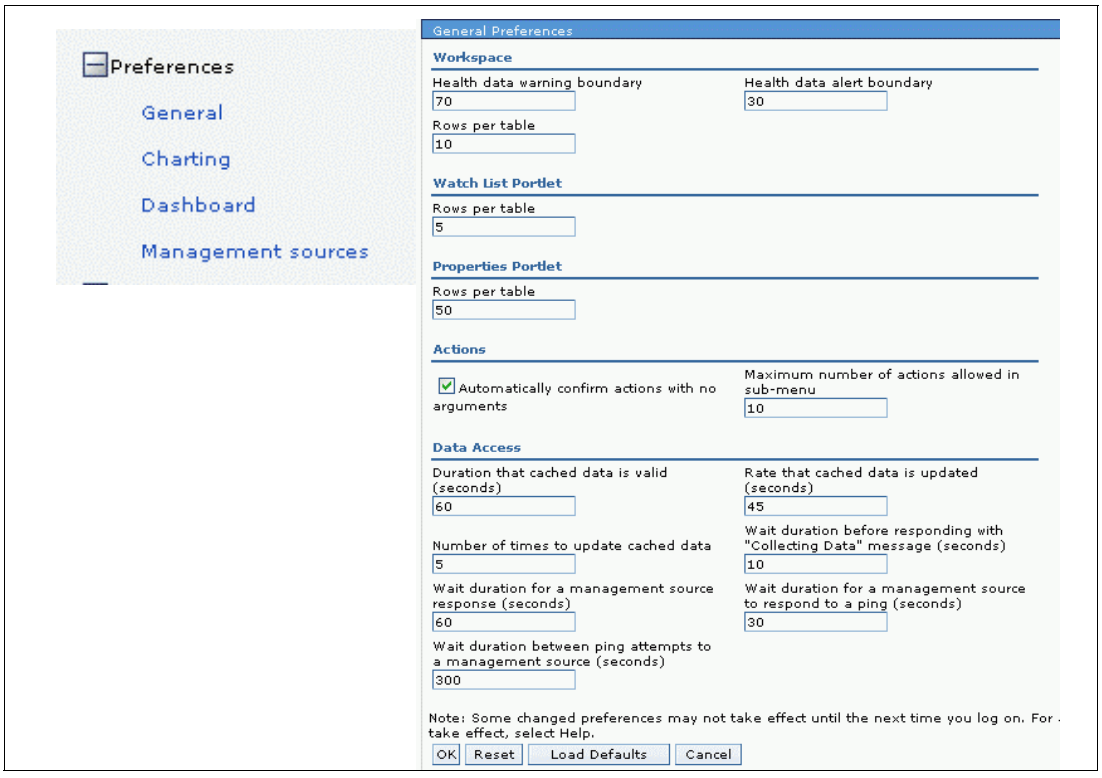

Figure 11-6 Virtualization Engine console General preferences

Virtualization Engine console uses LDAP server for its user administration. We can create and manage Virtualization Engine console users and groups by managing directory via Virtualization Engine console. The User and Group management menu is in the Settings tab. For example, if you want to add a new user in your administrator group, select **All portal user groups** → **iscadmins** (the administrator group for our environment) and then click **New user**, as shown in Figure 11-7. We add a group the same way that we add users.

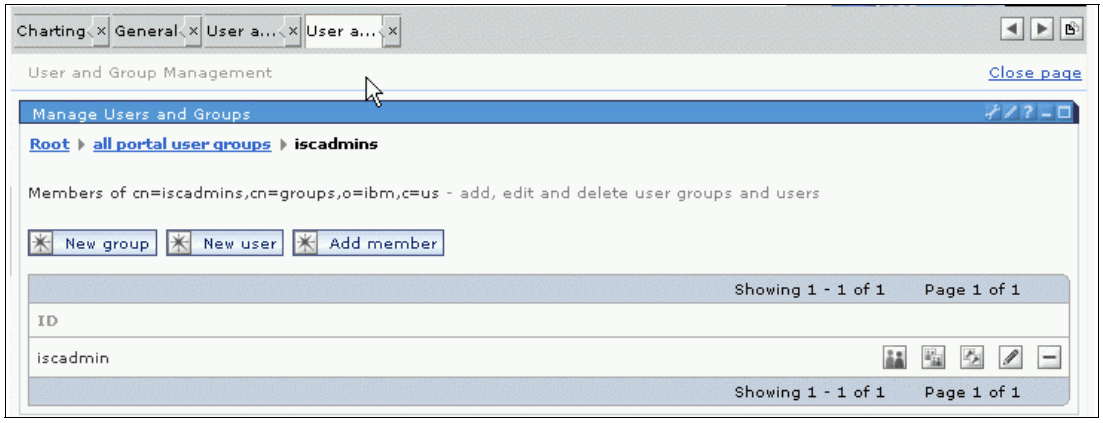

Figure 11-7 Adding a user into the Virtualization Engine console

Understand that ISC users and groups are not the same as the server users and groups. You can create users for the ISC using the ISC menus. However, these users are different from the users which have been created to access a server.

Management sources

For the Virtualization Engine Console to monitor the systems, it needs a management source definition. Before we use the Health Center, we need to add Management Source to the Virtualization Engine console.

You cannot install a management source from the Virtualization Engine console Health Center. Instead, you point to the location of a pre-configured and deployed Management Source to pass data from the Management Source to the Virtualization Engine console Bridge and then to the Virtualization Engine console Health Center.

To add Management Source, click **Management sources** in Preferences on the left navigation panel. Then, click **Add**. The configuration window opens. Select the source type and enter the parameters. Figure 11-8 shows an example.

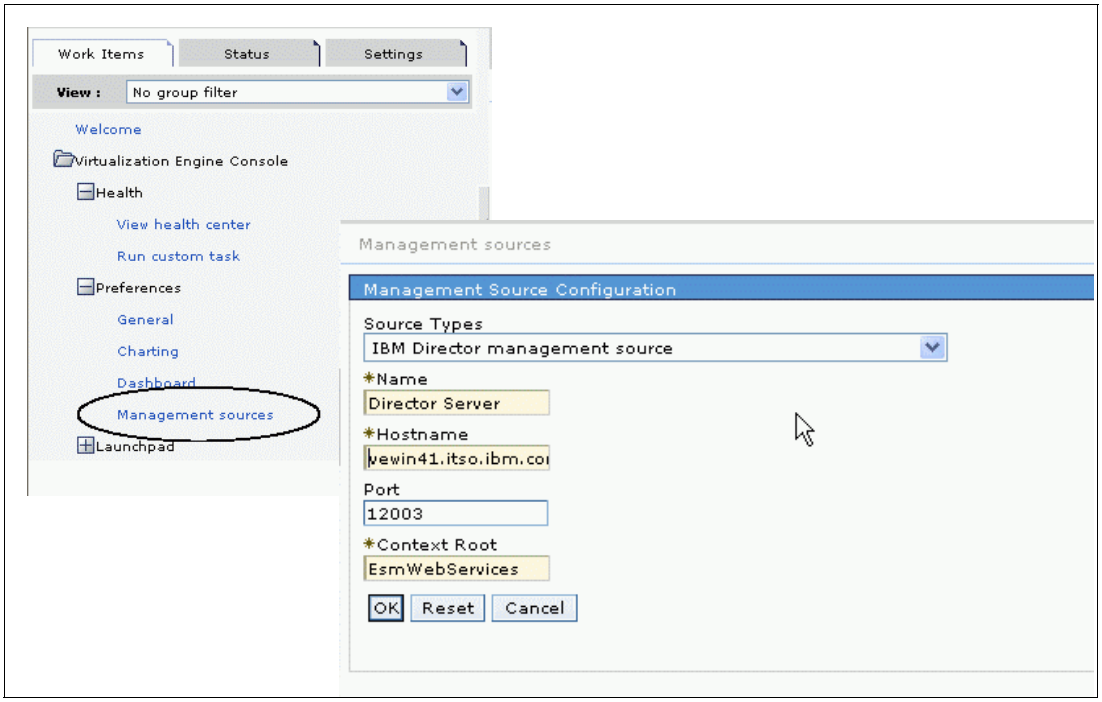

Figure 11-8 Configuring a management source

We use the IBM Director as our management source, and we used the threshold definitions that we created (as shown in Figure 11-5 on page 267) as input to the Virtualization Engine Console.

Under the Work Items tab in the Virtualization Engine Console, we select **Virtualization Engine Console → Health → View Health Center**. From the Health Center Summary section (at the right), select **Monitors** as shown in Figure 11-9.

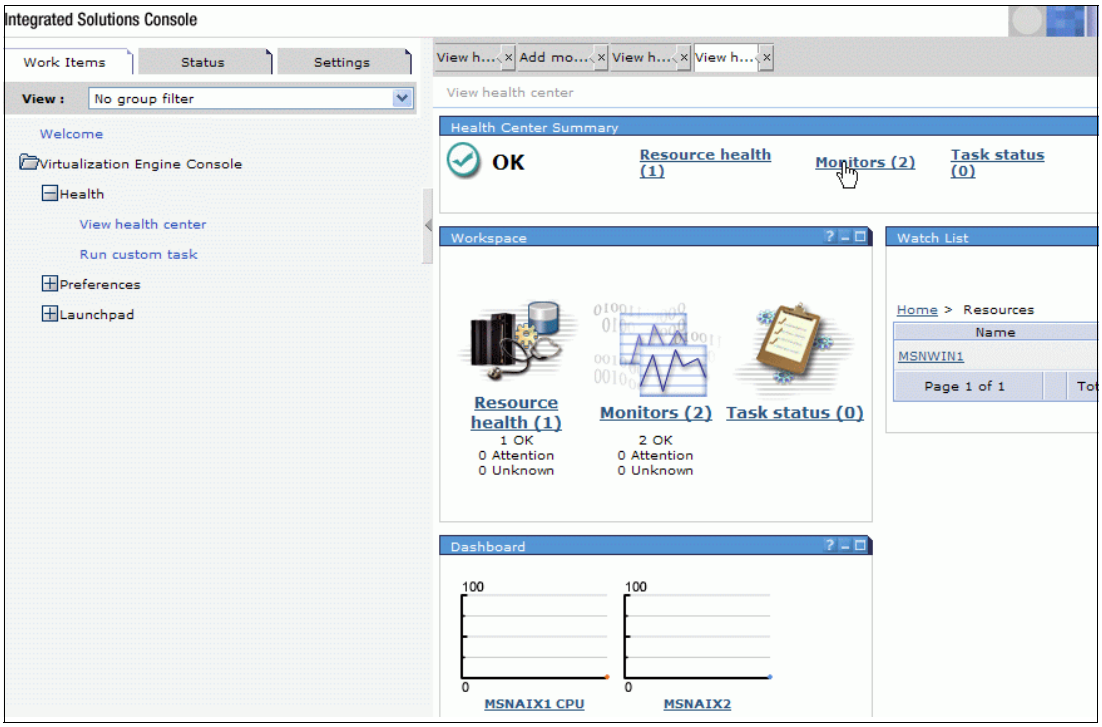

Figure 11-9 Select Monitors from the Virtualization Engine Console

We then select **Add a Monitor** to add the monitors that we just defined in the IBM Director. The available monitors from which we can select are shown in Figure 11-10.

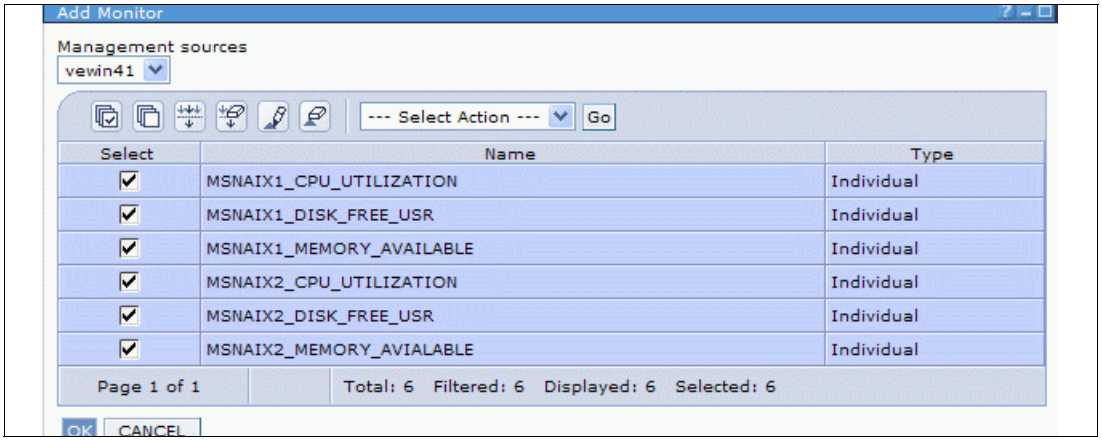

Figure 11-10 Add a monitor to Virtualization Engine console

The monitors that we defined in the IBM Director are presented automatically in the Virtualization Engine console to be selected as a resource to view from the Virtualization Engine Console. After we choose the monitors that we want to add, click **OK**. After returning to the Health Center, we can see that all the monitors were added successfully and are started, as shown in Figure 11-11.

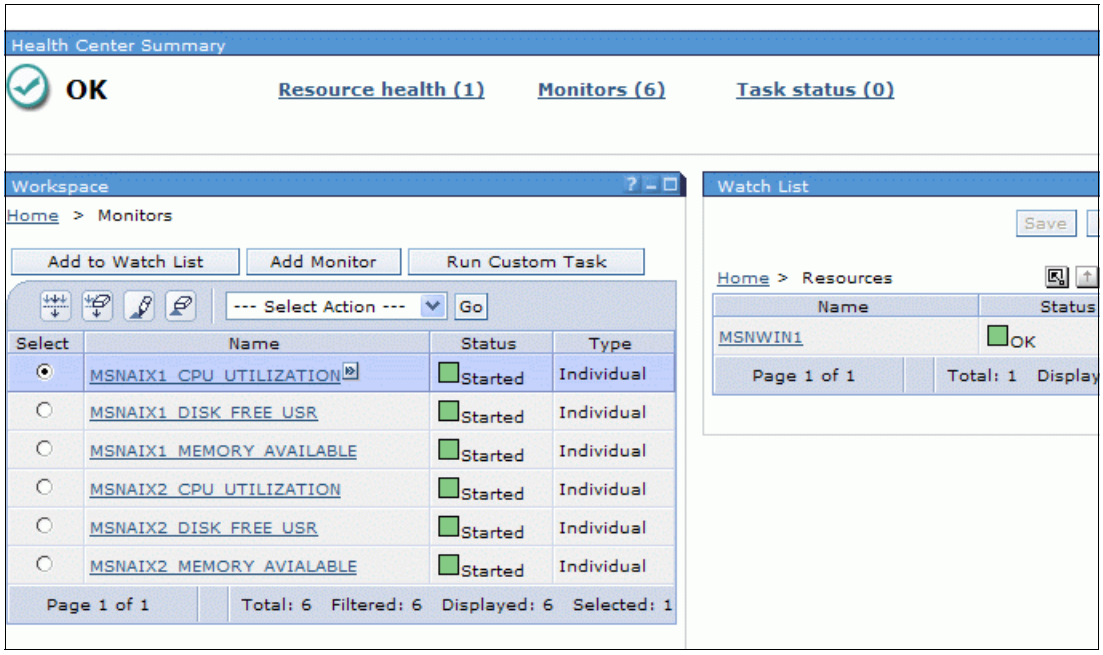

Figure 11-11 Monitors added to the Virtualization Engine *console*

If you want to stop a monitor, simply select the monitor, select the arrow next to the monitor (or click the pull-down), select **Monitor actions**, then select **Stop**. The status changes to grey and says stopped. You can always start it again by selecting the arrow next to the monitor, selecting **Monitor actions**, and then selecting **Start** as shown in Figure 11-12.

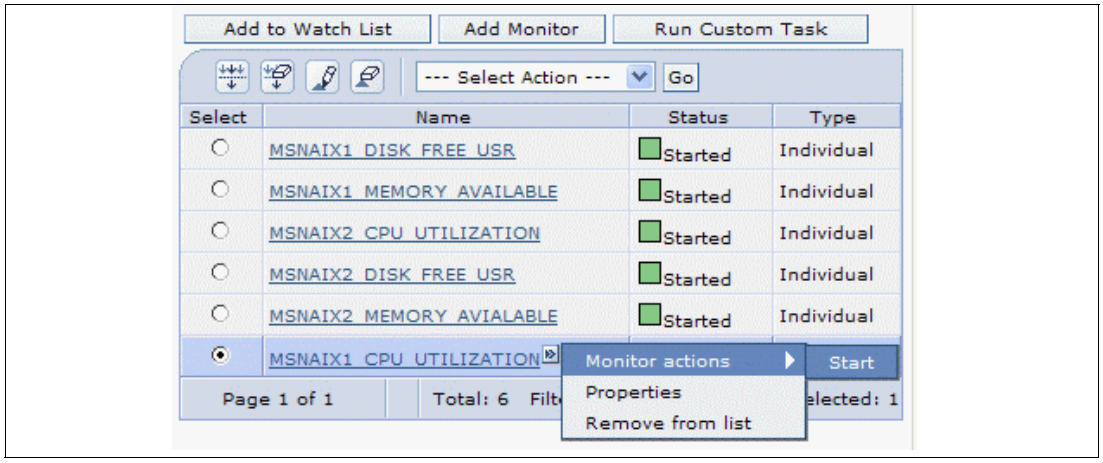

Figure 11-12 Virtualization Engine console - start a monitor

Depending on the resource that we are monitoring, we might want to add it to the watch list as we did for MSNWIN1 previously. We might want to add the monitor to the dashboard, as we did for MSNAIX1 CPU and MSNAIX2 CPU to have a running display of the monitor, such as the processor over time, as shown in Figure 11-13.

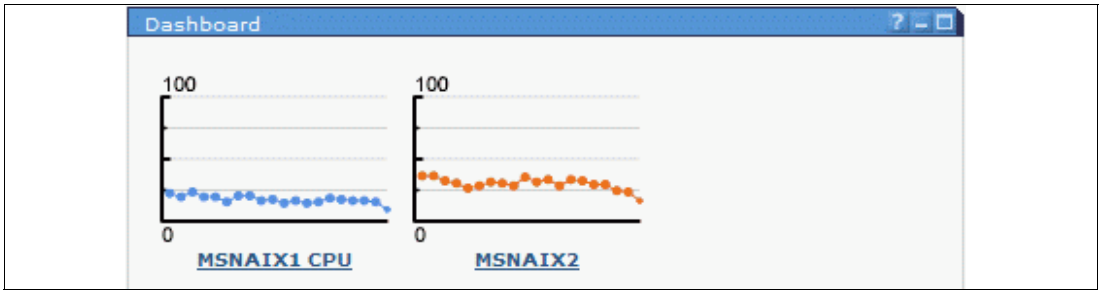

Figure 11-13 Virtualization Engine console Dashboard consisting of CPU for msnaix1 and msnaix2

We want to add some resources in our environment to monitor. We select **Resource health** from the Health Center summary and then select **Add resource**. We are presented with a list of available resources from the management source. We select several of the relevant resources, which are listed in Figure 11-14. After we choose the resources that we want to monitor in the Virtualization Engine Console, we select **Add resource** from the pull-down and click **OK**.

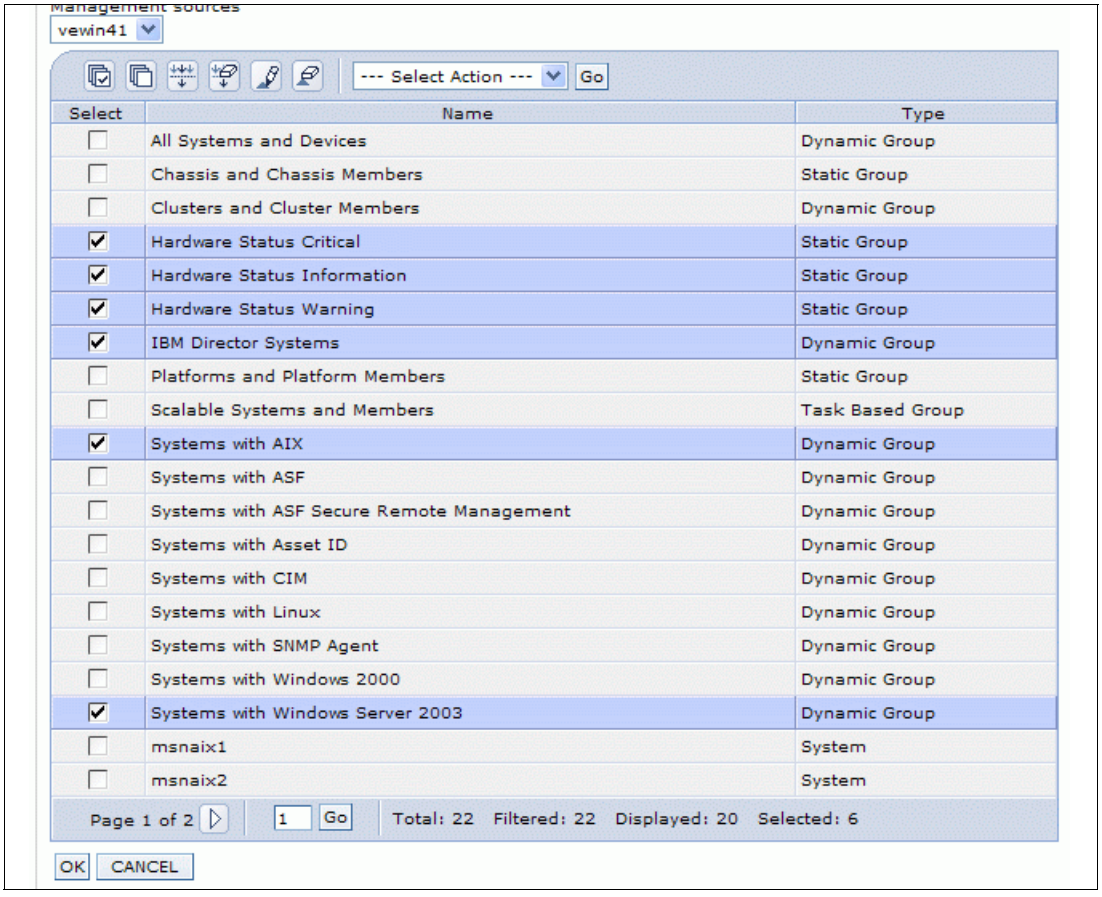

Figure 11-14 Resources selection for the Virtualization Engine console to monitor

After selecting the resource and clicking **OK**, we see they have the status OK as shown in Figure 11-15.

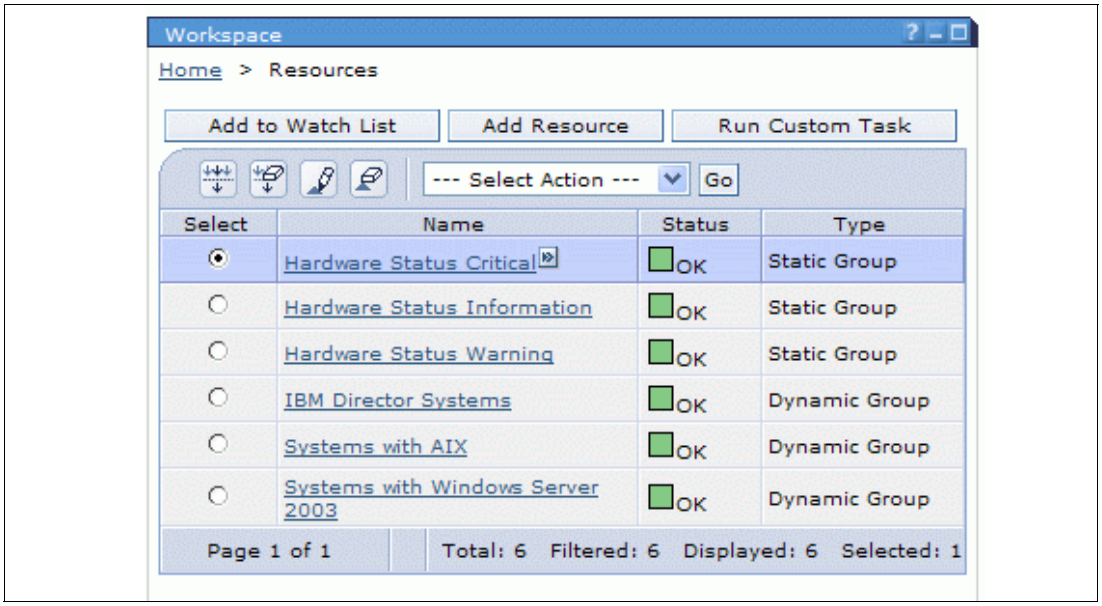

Figure 11-15 Resources selected

Just as we could do with the monitors that we defined previously, we can add any of these resources to the watch list by selecting the resource and clicking **Add to Watch List**.

Tip: You can change or pause the refreshing of the Health Center with the refresh portlet at the bottom of the Health Center (refer to Figure 3-11 on page 52).

Summary

For the Virtualization Engine Console configuration, we defined both monitors and resources that were defined in the IBM Director. We set up our HTTP server, MSNWIN1, to be part of the watch list and the processor of our two WebSphere Application Server servers, MSNAIX1 and MSNAIX2, to be included in the dashboard.

11.2.4 Tivoli Provisioning Manager consideration

We used Tivoli Provisioning Manager to provision resources into our environment. Tivoli Provisioning Manager can be configured to provision many different types of resources. For our heterogeneous scenario, we created a situation which would require a new WebSphere Application Server into our EWLM management domain to be provisioned, so we configured Tivoli Provisioning Manager to provision these resources for us. See 10.3.2, “Configuring Tivoli Provisioning Manager” on page 247 for a detailed description of the steps involved in created a Tivoli Provisioning Manager workflow to provision the EWLM Managed Server and the WebSphere Application Server in an AIX environment.

Now that the management servers have been configured, we need to configure the middleware products. Let’s discuss what is involved in this enablement.

11.2.5 WebSphere Application Server configuration

To run the Trade3 workload, we created a WebSphere Application Server cluster, called TradeCluster using the two AIX WebSphere Application Servers, shown in Figure 11-16. We installed the Trade3 enterprise application into the TradeCluster. This way, anytime a new WebSphere Application Server member is added to the cluster, the Trade3 application would be defined to it automatically.

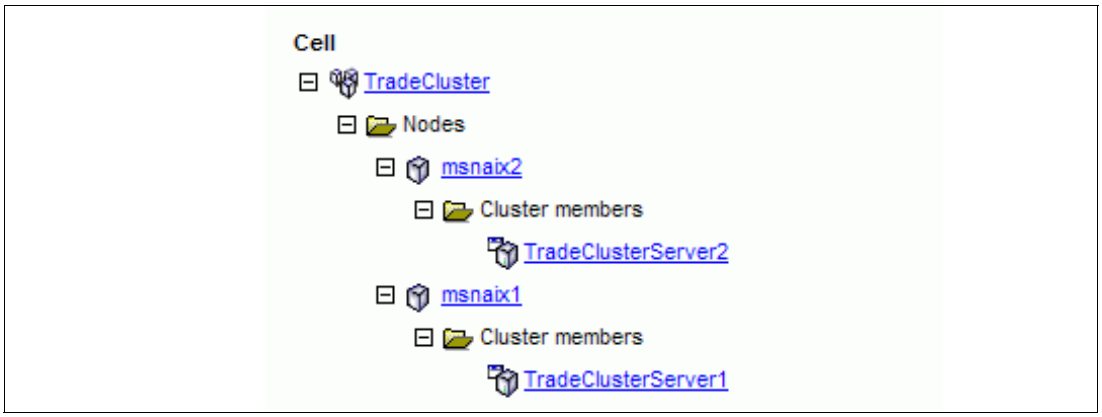

Figure 11-16 WebSphere Application Server Trade3 cluster

We needed to enable the WebSphere Application Server for ARM processing.

Note: Two WebSphere Application Servers were configured in advance, one to be used for the first part of the scenario, and one idle, for Tivoli Provisioning Manager provisioning later. We chose to preconfigure rather than have Tivoli Provisioning Manager create a second WebSphere application server because we were not testing Tivoli Provisioning Manager, just the provisioning aspect and how it works in our heterogeneous scenario

Enabling WebSphere Application Server for ARM

Perform the following steps to instrument WebSphere Application Server:

1. Enable PMI Request Metrics.
2. Create Java Virtual Machine custom properties.

Enable PMI Request Metrics

When Performance Monitoring Infrastructure (PMI) Request Metrics is enabled, WebSphere Application Server calls the ARM interface to pass the data to Enterprise Workload Manager. This information can be viewed through the EWLM Control Center and system log file. The correlation information is generated and logged for the Web server plug-in and the JDBC™ drivers.

To enable PMI Request Metrics:

1. Log in to the WebSphere Administrative Console.
2. Click **Troubleshooting** → **PMI Request Metrics**.
3. Enable Request Metrics and ARM by selecting both properties as shown in Figure 11-17 on page 276. Leave the Trace Level set to `HOPS`.

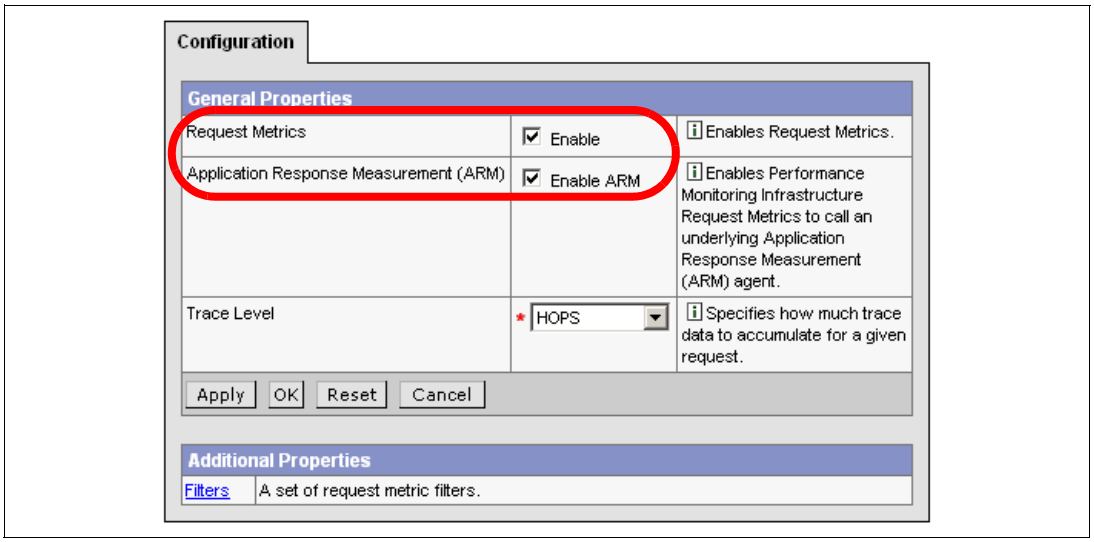

Figure 11-17 Enable Request Metrics and Application Response Management

4. Click **OK**.

Create Java Virtual Machine custom properties

When you enable PMI Request Metrics, you need to set the proper JVM custom properties for the type of ARM agent and the ARM transaction factory. You have to create the custom properties for all application servers you want Enterprise Workload Manager to monitor. Also, these settings affect the WebSphere HTTP plug-in configuration.

Create Java Virtual Machine custom properties using the following steps:

1. Log in to the WebSphere Administrative Console.
2. Click **Servers** → **Application Servers** → ***server name*** → **Process Definition** → **Java Virtual Machine** → **Custom Properties**.
3. If the properties shown in Table 11-1 are not present in the list, create them as shown in Figure 11-18 on page 277.

Table 11-1 Java Virtual Machine custom properties

Name	Value
ArmTransactionFactory	com.ibm.wlm.arm40SDK.transaction.Arm40 TransactionFactory
com.ibm.websphere.pmi.reqmetrics.ARMI MPL	arm4
com.ibm.websphere.pmi.reqmetrics.PassC orrelatorToDB	true
com.ibm.websphere.pmi.reqmetrics.logging Enabled	true
java.library.path	<EWLMMS_LIBPATH>
ws.ext.dirs	<EWLMMS_HOME>/classes

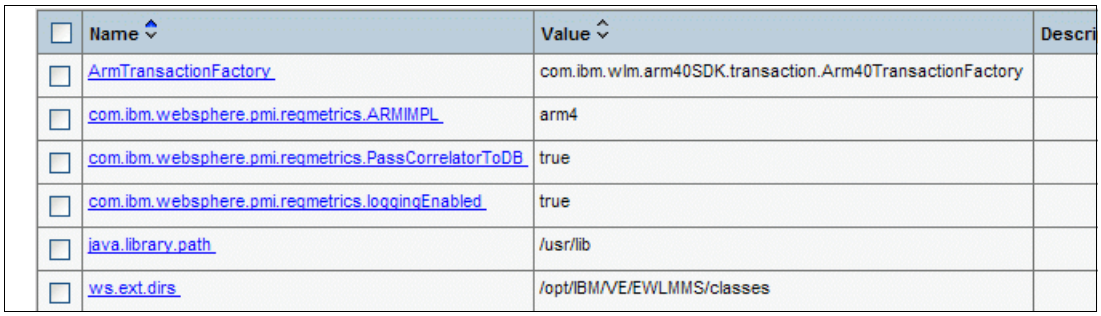

Name	Value	Descri
ArmTransactionFactory	com.ibm.wlm.arm40SDK.transaction.Arm40TransactionFactory	
com.ibm.websphere.pmi.reqmetrics.ARMIMPL	arm4	
com.ibm.websphere.pmi.reqmetrics.PassCorrelatorToDB	true	
com.ibm.websphere.pmi.reqmetrics.loggingEnabled	true	
java.library.path	/usr/lib	
ws.ext.dirs	/opt/IBM/VE/EWLMMS/classes	

Figure 11-18 Java Virtual Machine custom properties

Note: The value of com.ibm.websphere.pmi.reqmetrics.loggingEnabled is an optional property. If you set this parameter to `true` and the PMI Request Metrics trace level is set to `HOPS`, `PERF_DEBUG`, or `DEBUG`, you see the correlator information in the application server's SystemOut.log as shown in Example 11-1. Note that if this property is not created, the correlator information is not logged.

Example 11-1 PMI request metrics log at SystelOut.log

```
[6/7/05 11:30:57:285 CDT] 3f3bd9f5 PmiRmArmWrapp I PMRM0003I:  parent:ver=1,ip=1
0.1.1.190,time=1118084316874,pid=356532,reqid=4780365,event=1 - current:ver=1,ip
=10.1.1.190,time=1118084316874,pid=356532,reqid=4780378,event=1 type=JDBC detail
=SELECT T1.LOW, T1.OPEN1, T1.VOLUME, T1.PRICE, T1.HIGH, T1.COMPANYNAME, T1.SYMBO
L, T1.CHANGE1 FROM QUOTEEJB  T1 WHERE T1.SYMBOL = ? elapsed=5
[6/7/05 11:30:57:323 CDT] 3f2bd9f5 PmiRmArmWrapp I PMRM0003I:  parent:ver=1,ip=1
0.1.1.184,time=1117714471390,pid=2600,reqid=699966,event=1 - current:ver=1,ip=10
.1.1.190,time=1118084316874,pid=356532,reqid=4788544,event=1 type=URI detail=/tr
ade/app elapsed=1
[6/7/05 11:30:57:323 CDT] 3f3bd9f5 PmiRmArmWrapp I PMRM0003I:  parent:ver=1,ip=1
0.1.1.190,time=1118084316874,pid=356532,reqid=4780365,event=1 - current:ver=1,ip
=10.1.1.190,time=1118084316874,pid=356532,reqid=4780379,event=1 type=JDBC detail
=SELECT T1.LOW, T1.OPEN1, T1.VOLUME, T1.PRICE, T1.HIGH, T1.COMPANYNAME, T1.SYMBO
L, T1.CHANGE1 FROM QUOTEEJB  T1 WHERE T1.SYMBOL = ? elapsed=2
```

4. Click **Save**.

   For the configuration changes to take effect, the changes must be saved to the master configuration and synchronized to the nodes.

5. Check **Synchronize changes with nodes** and click **Save** to save all your changes.
6. Restart the application server.

Note: If a WebSphere Application Server Network Deployment is present in your installation, set these properties for the deployment manager as well or you might see two problems:

1. The HTTP plug-in configuration regeneration cannot set the correct loggingEnable value in the plug-in configuration file (plugin-cfg.xml).
2. Future releases might not be able to pick up the property.

On the other hand, if you do not want to instrument the Network Deployment, you can edit the plugin-cfg.xml for `loggingEnable=false`.

Enabling IBM HTTP Server for ARM

There are two instrumented plug-ins for IBM HTTP Server:

- The WebSphere HTTP plug-in that is provided by WebSphere installation
- The independent plug-in that is provided for nonWebSphere Application Server application usage

Also, there are two versions of IBM HTTP Server that are supported by Enterprise Workload Manager:

- IBM HTTP Server 1.3.28; only the WebSphere HTTP plug-in is available for this version.
- IBM HTTP Server 2.0.47; both the WebSphere HTTP plug-in and the independent plug-in are available.

Instrumenting IBM HTTP Server depends on which version of IBM HTTP Server you are currently running and what type of plug-in you use. We ran with the IBM HTTP server 1.3.28 and the WebSphere HTTP plug-in.

WebSphere HTTP plug-in

The WebSphere HTTP plug-in receives a user's HTTP request and determines whether the request should be handled by the Web server or an application server using the plug-in configuration file, plugin-cfg.xml. If a request for an application server is received, the plug-in dispatches it to the appropriate Web application with ARM correlator information.

The WebSphere HTTP plug-in that ships with WebSphere Application Server is already instrumented, so all you have to do is install the plug-in to the IBM HTTP Server and update the plug-in configuration file after WebSphere Application Server instrumentation is successfully completed.

Perform the following steps to instrument WebSphere HTTP plug-in:

1. Install the WebSphere HTTP plug-in to IBM HTTP Server using the install wizard.
2. Apply WebSphere Application Server Fix Pack 1, if needed.
3. Verify the IBM HTTP Server configuration file.

   For IBM HTTP Server 1.3.28, add the following line at <IHS13_HOME>/conf/httpd.conf:

   ```
   LoadModule ibm_app_server_http_module <WAS_HOME>/bin/mod_ibm_app_server_http.so
   WebSpherePluginConfig <WAS_HOME>/config/cells/plugin-cfg.xml
   ```

   For IBM HTTP Server 2.0.47, add the following line at <IHS20_HOME>/conf/httpd.conf:

   ```
   LoadModule was_ap20_module <WAS_HOME>/bin/mod_ibm_app_server_http.so
   WebSpherePluginConfig <WAS_HOME>/config/cells/plugin-cfg.xml
   ```

4. Update the plug-in configuration file.
   a. Log in to the WebSphere Administrative Console.
   b. Click **Environment** → **Update Web Server Plug-in**.
   c. Click **OK** to regenerate the Web server plug-in configuration file.
   d. Open the plug-in configuration file <WAS_HOME>/config/cells/plugin-cfg.xml and search for this line:

      ```
      <RequestMetrics armEnabled="true" loggingEnabled="false" rmEnabled="true"
      traceLevel="HOPS">
      ```

      Ensure that loggingEnabled and traceLevel attributes match the values you specified at PMI Request Metrics described previously. If the loggingEnabled parameter is `true`

and **traceLevel is s**et to other than NONE, you see the correlator information in HTTP plug-in log file as shown in Example 11-2.

Example 11-2 The correlator information

```
[WedJun 6 Jun 6 11:28:31 2005] 00000658 000008b4 - PLUGIN:
parent:ver=1,ip=9.12.4.1849.12.4.184,time=1085613001328,pid=1624,reqid=49,event=1 -
current:ver=1,ip=9.12.4.184,time=1085613001328,pid=1624,reqid=49,event=1 type=HTTP
detail=/PlantsByWebSphere/servlet/ShoppingServlet elapsed=16 bytesIn=0 bytesOut=12520
[Wed Jun 6 11:28:31 2005] 00000658 000008c0 - PLUGIN:
parent:ver=1,ip=9.12.4.184,time=1085613001328,pid=1624,reqid=50,event=1 -
current:ver=1,ip=9.12.4.184,time=1085613001328,pid=1624,reqid=50,event=1 type=HTTP
detail=/trade/app elapsed=0 bytesIn=0 bytesOut=3057
[Wed Jun 6 11:28:31 2005 11:28:31 2005] 00000658 000008bc - PLUGIN:
parent:ver=1,ip=9.12.4.184,time=1085613001328,pid=1624,reqid=51,event=1 -
current:ver=1,ip=9.12.4.184,time=1085613001328,pid=1624,reqid=51,event=1 type=HTTP
detail=/PlantsByWebSphere/servlet/ShoppingServlet elapsed=16 bytesIn=0 bytesOut=12520
```

e. The plug-in configuration file plugin-cfg.xml, must reside in the location that is specified in the IBM HTTP Server configuration file. Open the IBM HTTP Server configuration file stored in <IHS13_HOME>/config/httpd.conf or <IHS20_HOME>/config/httpd.conf and search for the line:

```
WebSpherePluginConfig "<WAS_HOME>/config/cells/plugin-cfg.xml"
```

Transfer the updated plugin-cfg.xml file to this path on the IBM HTTP Server machine.

5. Restart the IBM HTTP Server.

11.2.6 DB2 consideration

Currently, DB2 V8.2 and WebSphere Application Server V6.0.1 are the available middleware products available for z/OS that have the ARM enablement capability. The following maintenance is required to provide ARM enablement on DB2 Universal Database V8.2. When you apply the following maintenance, DB2 is ARM enabled by default and there are no additional steps required to enable ARM on DB2 for z/OS managed servers.

- APAR PQ91509 EWLM correlator support to allow to do performance measurement using ARM instrumentation
- APAR PQ99707 with the connection reset optimization code for JCC clients.
- The JCC driver should be at 2.6.16 level or above.

Transactions initiated through DDF are accounted as separate transactions. DB2 transactions initiated by a local WebSphere Application Server are executed in binding mode and so forth accounted to the WebSphere thread.

In this release, DDF does not create a correlator for the transactions but it will only handle the correlator coming from a distributed environment.

For more information, refer to *Enterprise Workload Manager z/OS Support*, REDP-3993.

11.3 Validating the configuration

We validated our configuration with the EWLM control center and the Virtualization Engine Console.

11.3.1 Validating the configuration with EWLM

We used the following displays from EWLM:

- We used the EWLM control center managed server view to validate that all servers were part of the EWLM management domain by selecting **Manage** → **Managed servers** and verifying that the state of all of the systems were active as shown in Figure 11-19. Note, it also displays the operating system and the release level it is running.

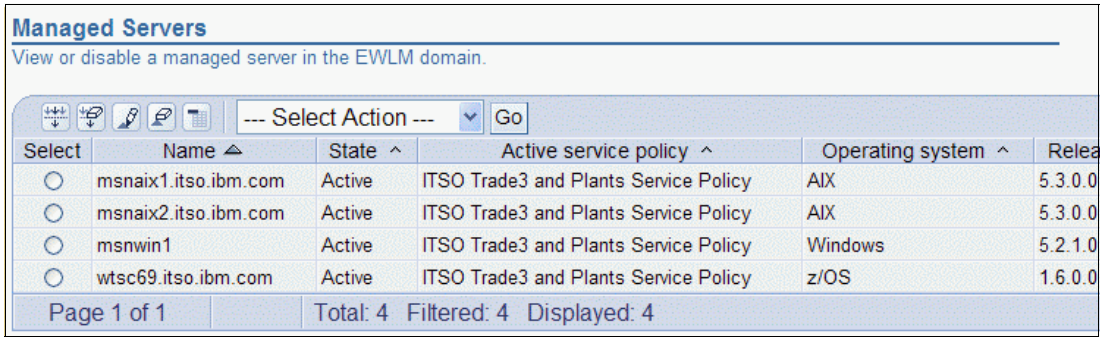

Figure 11-19 Verification of the EWLM Managed Servers connected to the EWLM DM

- We validated that the transactions were being reported correctly in the correct transaction class and service class my executing the transactions manually and verifying the Service classes and Transaction classes monitor views were showing the expected data.
- We verified that the flow of the transactions from HTTP to WebSphere to DB2 was correct by selecting one of the Trade3 service classes, and selecting the pull-down option Application topology as shown in Figure 11-20.

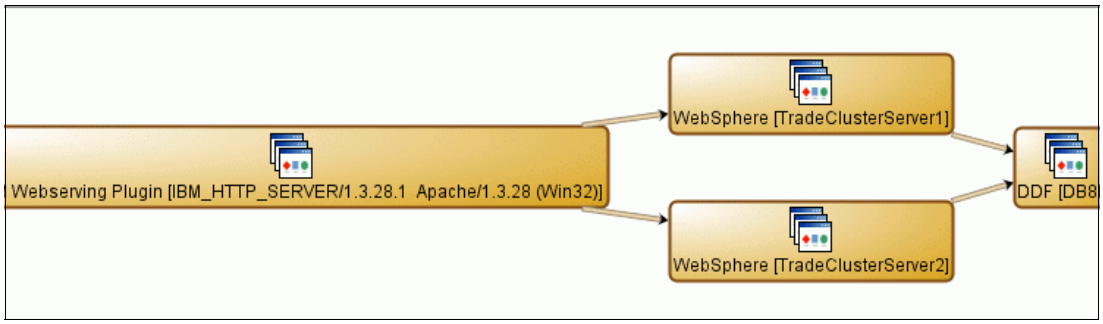

Figure 11-20 Verifying the application topology view

11.3.2 Validating IBM Director and the Virtualization Engine console

We validated that the Windows and AIX servers were part of the IBM Director management domain and reporting in the Virtualization Engine Console by viewing the Virtualization Engine Console Health Center and selecting **Monitors.**

The view displayed the monitors in the workspace and the resources under the watch list. All monitored resources were started and the status was OK (and denoted by a green box next to the word *Started* or *OK*), as shown in Figure 11-21.

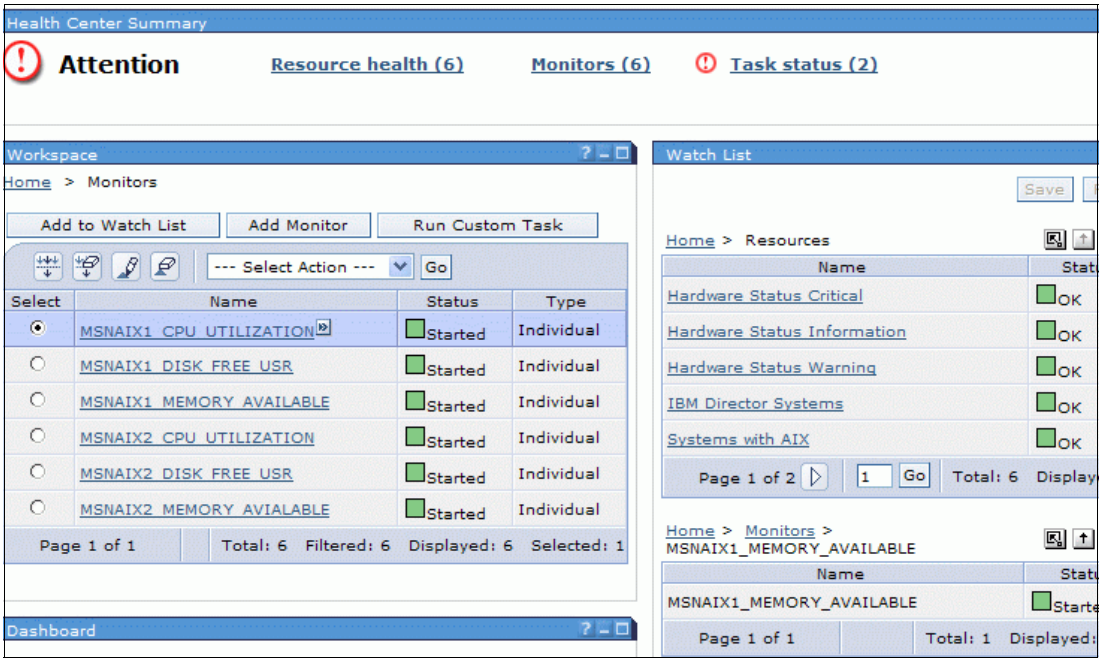

Figure 11-21 Validation of Virtualization Engine console defined monitors and resources

11.4 Testing the heterogeneous scenario

Now that the setup has been validated, we wanted to start our test. We stopped the second WebSphere Application Server running on MSNAIX1 and stopped the EWLM Managed Server. We received a communication error message in the EWLM control center in the **Monitor** → **Managed servers** view, as expected. At this point we were ready to begin the first part of our scenario.

11.4.1 Monitoring the workload running under normal conditions

This part of the scenario uses the EWLM control center and the Virtualization Engine Console to monitor the environment running under *normal* conditions. It starts with the EWLM performance goals being met and all of the systems running without problems or exceptions. In this test, the workload is running with one HTTP Server, one WebSphere Application Server and one database manager, DB2.

EWLM control center

The EWLM control center was the focal part for us to verify if our performance goals are being met. We began looking at the service classes to verify that the PI was being met, as shown in Figure 11-22.

Service class	PI	Importance	Performance	Goal
SC_Calc_batch	0.00	Medium	Slowest velocity	Slowest velocity
SC_InternetExplorer	0.00	Medium	Slowest velocity	Fast velocity
SC_Trade3_account	0.80	Medium	00.048 average	00.060 average
SC_Trade3_default	0.80	Medium	00.031 average	00.040 average
SC_Trade3_portfolio	0.41	Medium	00.032 average	00.080 average
SC_Trade3_profile	0.80	Medium	00.079 average	00.100 average
SC_Trade3_quotes	0.44	Medium	00.013 average	00.030 average
SC_Trade3_shopping	0.87	Medium	94% in 00.300	90% in 00.300

Figure 11-22 Service classes running under normal non-stressed conditions

The EWLM control center also allows us to view the response time which we are achieving over time. We looked at the view for the SC_Trade3_portfolio as shown in Figure 11-23.

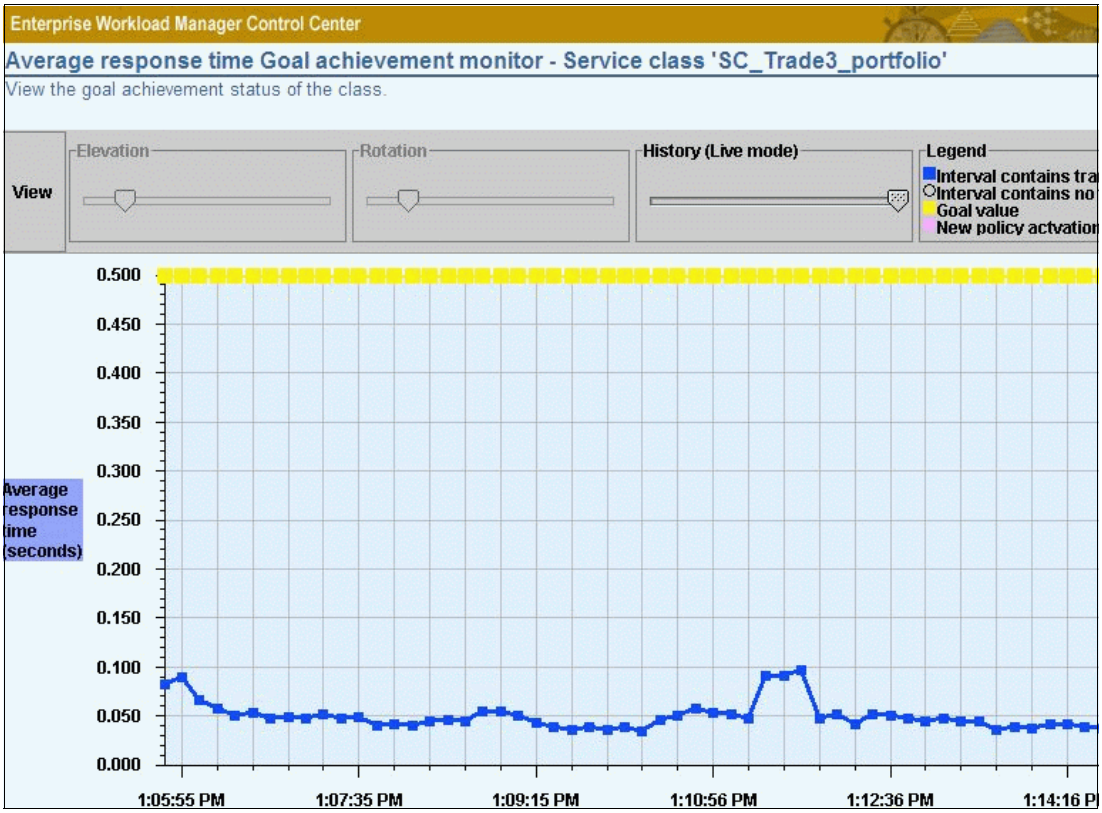

Figure 11-23 Response time Goal achievement monitor

After we then verified that all of the PIs were being met, we wanted to check the system management thresholds which we set to verify that everything was running smoothly.

Virtualization Engine Console

In the Virtualization Engine Console, we verified that there were no monitors or resources being flagged. We also looked at the processor to see if they if there were any trends or reaching any type of warning level. We verified that everything was still running well and that there were no problems as seen in Figure 11-13 on page 273.

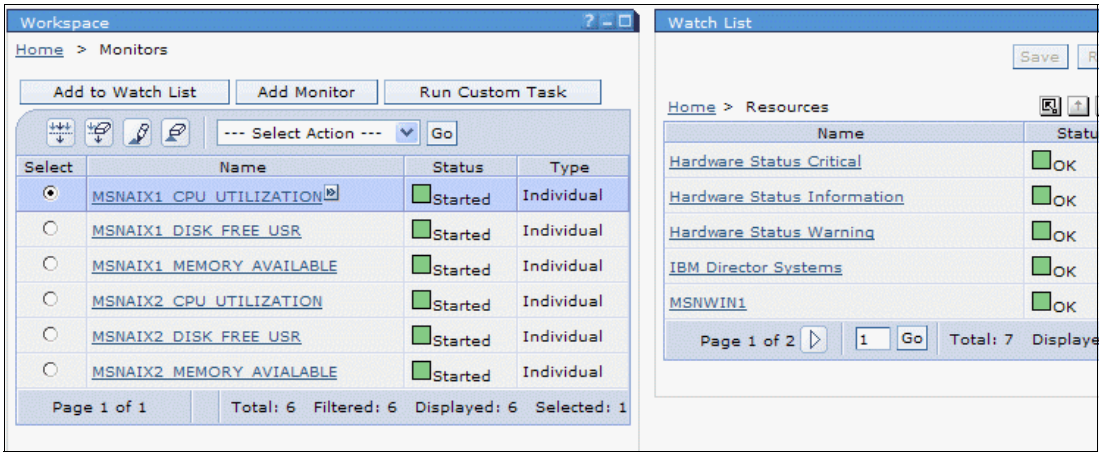

Figure 11-24 Virtualization Engine console everything OK

If anything was flagged as yellow, this would indicate that a low warning threshold had been met and would require further investigation. Because all of our thresholds were within our expected range, we did not need to investigate further. We were able to continue with part two of our scenario, putting the system under stressful conditions.

11.4.2 Monitoring the workload under stressful condition

In this part of our scenario, we ran some extra work on the AIX machine to drive the CPU higher which would cause the WebSphere Application Server to contend for CPU resources. We used the monitoring tools to verify the effects of adding extraneous work to the msnaix2 system which drove the system harder.

EWLM control center

We began seeing that the EWLM PIs were not met by monitoring the service classes in the EWLM control center as shown in Figure 11-25. We got this view by selecting **Monitor** → **Service classes**. The four PIs which exceed 1.0 also show up in the **Monitor** → **Exceptions** view.

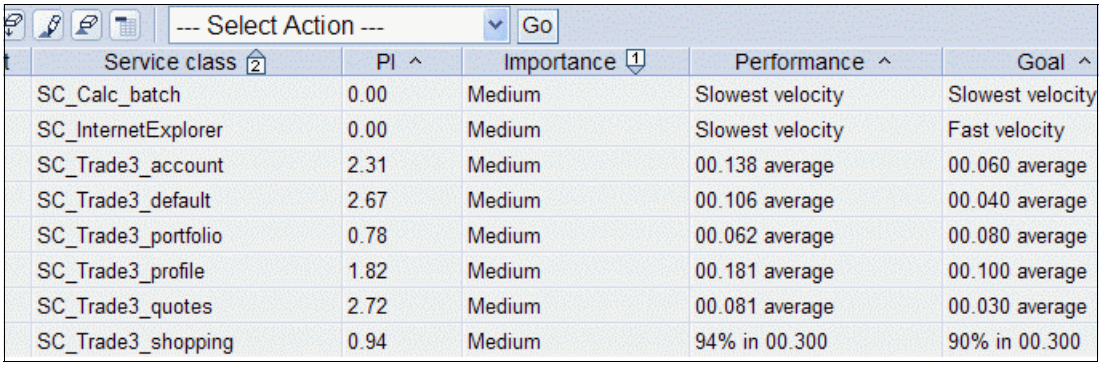
--- Select Action ---
Go
Service class
PI
Importance
Performance
Goal
SC_Calc_batch 0.00 Medium Slowest velocity Slowest velocity
SC_InternetExplorer 0.00 Medium Slowest velocity Fast velocity
SC_Trade3_account 2.31 Medium 00.138 average 00.060 average
SC_Trade3_default 2.67 Medium 00.106 average 00.040 average
SC_Trade3_portfolio 0.78 Medium 00.062 average 00.080 average
SC_Trade3_profile 1.82 Medium 00.181 average 00.100 average
SC_Trade3_quotes 2.72 Medium 00.081 average 00.030 average
SC_Trade3_shopping 0.94 Medium 94% in 00.300 90% in 00.300

Figure 11-25 Verifying the stressed system with EWLM control center

We selected the application topology view for SC_Trade3_default as that was the service class which had exceeded its response time goal the most. From there we selected the **tableview** and reviewed the data. We noticed that there was a Processor delay in the WebSphere Application Server, however the amount of processor that the was using was rather low, as shown in Figure 11-26. We determined that it might be caused by some of the other activity running on the system. We verified the other service classes which had exceptions had the same symptoms as SC_Trade3_default by looking that the same views for those service classes.

Hop	Name	Type	Platform	Average response time...	Average active time (se...	Processor time (secon...	Pr
0	IBM Webserving Plugin [IBM_HTTP...	Application		00.106	00.001	00.0	0%
0	msnwin1	Server	Windows	00.106	00.001	00.0	0%
0	MSNWIN1/PID=0000002600	Instance		00.106	00.001	00.0	0%
1	WebSphere [TradeClusterServer2]	Application		00.098	00.084	01.619720	18%
1	msnaix2.itso.ibm.com	Server	AIX	00.098	00.084	01.619720	18%
1	msnaix2.TradeClusterServer2	Instance		00.098	00.084	01.619720	18%
2	DDF [DB8E]	Application		00.001	00.001	00.126214	0%
2	wtsc69.itso.ibm.com	Server	z/OS	00.001	00.001	00.126214	0%
2	DB8E	Instance		00.001	00.001	00.126214	0%

Figure 11-26 EWLM tableview processor delay

Virtualization Engine Console

What we noticed is that the total processor use of AIX2 was starting to reach a threshold and showed up as a warning in the Virtualization Engine Console, as shown in Figure 11-27.

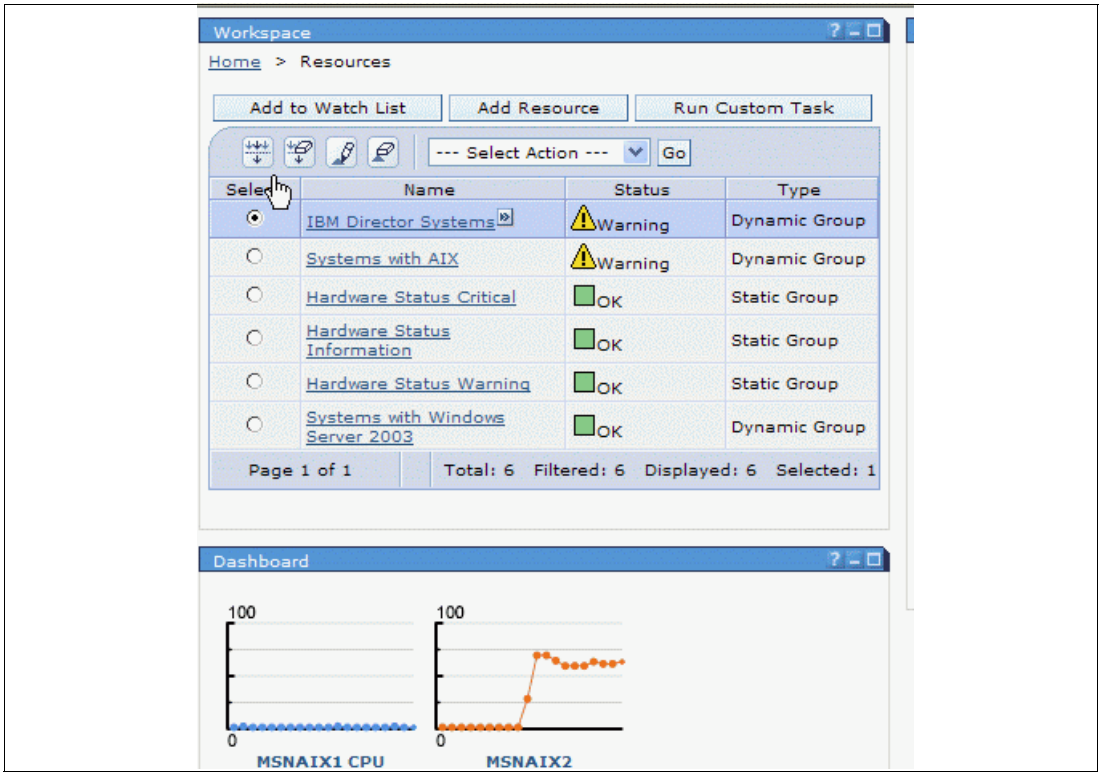

Figure 11-27 Virtualization Engine console warning threshold met

As the workload kept running, the processor for AIX2 went from a warning threshold to an alert as shown in Figure 11-28 because it hit another threshold.

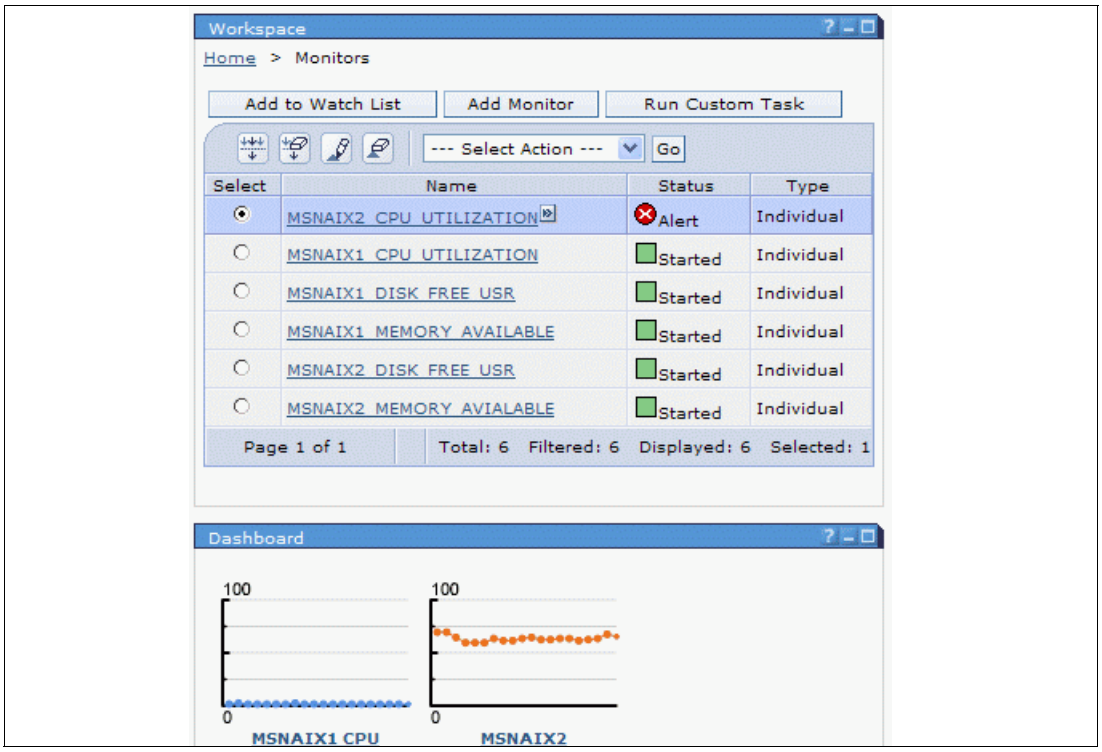

Figure 11-28 Virtualization Engine console alert

Because we did not define all of the tasks in the IBM Director, we used the IBM Director to analyses what processes were running on the system which were consuming most of the processor, which is under the Process Management function (similar to a **`ps -ef`** command, or on Windows. We determined that it was another process which was unrelated to our Trade3 workload. This could also have been performed using the task status function of the Virtualization Engine Console.

At this point we had a few choices to in order to maintain our PI levels:

- Stop (or move) the other job running on the system.
- Use the AIX workload manager function to limit the resources that the other job is consuming.
- Leave the other job running as is while provisioning another WebSphere Application Server to run on another system to absorb the overflow of work.

We decided to use Tivoli Provisioning Manager to provision a new server to offset the workload.

11.4.3 Resolving and monitoring the stressful conditions

At this point, our PIs were continuing to increase due to queuing and we needed to provision a new server to handle the Trade3 workload to meet the performance goals that we set previously.

Tivoli Provisioning Manager

We used Tivoli Provisioning Manager to do the following:

- Provision a second AIX into our EWLM domain management.
- Provision a second WebSphere Application Server on the second AIX system.

Because the HTTP which we were using was preconfigured with the definitions of the second WebSphere Application Server on msnaix1, we did not have to configure anything extra for the second WebSphere Application Server.

In particular, keep in mind the following:

- WebSphere Application Server provides plug-ins for several Web server brands and versions. Each Web server has specific tuning parameters that affect the application performance. One of these variables is called the RetryInterval. The RetryInterval is the length of time to wait before trying to connect to a server that has been marked temporarily unavailable. In our scenario, the WebSphere Application Server running on msnaix1 was temporarily unavailable until Tivoli Provisioning Manager provisioned it into our cluster. We needed to wait until the retry interval had been reached in the Web server plug-in in order for the HTTP server to retry the second WebSphere Application Server which was provisioned into the cluster (defined in the http plug-in, shown in Example 11-3 as RetryInterval=”60”, which is the default).

Example 11-3 Example of the http plug-in RetryInterval setting

```
<ServerCluster CloneSeparatorChange="false" LoadBalance="Round Robin" Name="TradeCluster"
PostSizeLimit="-1" RemoveSpecialHeaders="true" RetryInterval="60">
```

- Additionally, there is HTTP and WebSphere Application Server caching which occurs when the beginning transactions start executing in the HTTP and WebSphere Application Server, thus the response times will be higher until the WebSphere Application Server information has been cached.

EWLM control center

We continued to use the EWLM control center to monitor the effects of the provisioned system as follows:

- We validated that the second AIX system was part of the EWLM management domain by selecting **Manage** → **Managed servers** and verifying that the state of all of msnaix1 went from communication error to active.
- We validated the application topology by selecting **Application topology** for the Trade3 service classes. We verified that both of the WebSphere Application Servers were being used as shown in Figure 11-20 on page 280.
- We monitored the PI levels of the service classes and verified that the service class exceptions were becoming less. Again, we wanted to validate that the performance goal defined were being met.

Specific considerations include the following:

- Because there is a slight delay before the managed server sends the performance information to the domain manager (every 10 seconds) and another delay before the EWLM control center accumulates this new performance information and displays it in the views (30 seconds), we waited a few minutes before determining if our provisioning resolved our performance problem.
- Also, there is a preference setting within the EWLM control center to allow you to view the latest EWLM performance data from one minute to one hour. We set our interval to 5 minutes. Before making our determination, we waited until a full ten minutes had passed to view the performance information of the last five minutes so that we did not include the data from the previous part of the scenario. Over a period of 10 minutes or so, a 5 minute duration for our service class appeared to be running fine, as shown in Figure 11-29.

Select	Service class	PI	Importance	Performance	Goal
○	SC_Calc_batch	0.00	Medium	Slowest velocity	Slowest velocity
○	SC_InternetExplorer	0.00	Medium	Slowest velocity	Fast velocity
○	SC_Trade3_account	0.80	Medium	00.048 average	00.060 average
○	SC_Trade3_default	0.80	Medium	00.031 average	00.040 average
○	SC_Trade3_portfolio	0.41	Medium	00.032 average	00.080 average
○	SC_Trade3_profile	0.80	Medium	00.079 average	00.100 average
○	SC_Trade3_quotes	0.44	Medium	00.013 average	00.030 average
○	SC_Trade3_shopping	0.87	Medium	94% in 00.300	90% in 00.300

Figure 11-29 Service class back to normal Virtualization Engine Console

Virtualization Engine Console

When the second AIX was provisioned with the second WebSphere Application Server, we looked at the processor of both MSNAIX2 and the newly provisioned MSNAIX1. We saw that the work was beginning to get off loaded to MSNAIX1 and that the processor threshold of MSNAIX2 was back within a normal range (it dropped below the warning threshold), as shown in Figure 11-30.

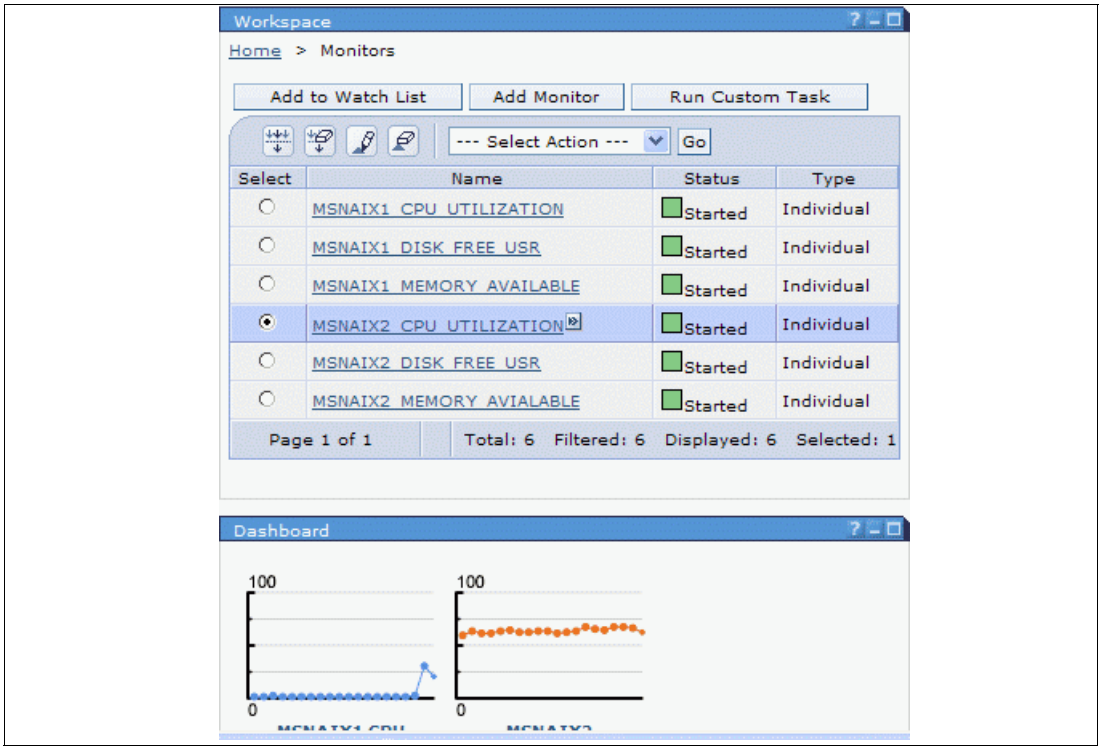

Figure 11-30 Work beginning to get off-loaded to MSNAIX1

11.5 Conclusions

In testing this scenario, we made the following conclusions:

- We found that both the Enterprise Workload Manager control center and the Virtualization Engine Console offered us a consolidated view in an non homogeneous environment which we previously didnt have before. When a workload starts to expand beyond a single system, and beyond a single operating system, having a product which spans these boundaries becomes more important. Having an up to date topology of where the workload traverses is of utmost importance. We can now tell if we want to do maintenance on a middleware product or an operating system, if a system is running under constraint, or if there are network related problems, which workload would be affected.
- If the help desk starts to get calls from customers about a slow down in the system, they now have a better place to start to reduce diagnostics time. By looking at the priority of the problem from an EWLM point of view, they can then determine how this problem relates to the business goals, and put the appropriate resources on the problem.
- Utilizing the views in the EWLM control center allowed us to not only monitor our systems, but the actual workload which was running on them from a Service Level Agreement point of view. We were able to determine if the service level agreements were being met, and if not, by how much. We were able to start to perform the first steps in diagnostics to help isolate the problem. Which could have taken us days to deduce where the problem

existed, we were able to narrow down the area of concern within minutes. Depending on the problem, and the priority (or the importance level) of the work, we could then decide if we should provision more resources to it, delay the work to a later time when resources were not as constrained or to leave it alone temporarily.

- Because we preconfigured the pool of provisioned resources into our existing environment, we were able to dynamically provisioned the new resource very efficiently. When the problem occurred due to lack of resources in the WebSphere Application Server, having the ability to limit the amount of Tivoli Provisioning Manager workflow definitions (plug-in update, application server and enterprise application definitions), and Tivoli Provisioning Manager workflow execution simplified our scenario. If we did not have the ability to preconfigure these resources, we would have been able to use Tivoli Provisioning Manager to provision them for us, but then it would have taken more time, which in turn would have made the turnaround time for the problem resolution to be extended.
- We were able to preconfigure the HTTP server with the definition of the second WebSphere Application Server and to preconfigure the second Application Server which ran on MSNAIX1 into the WebSphere cluster, TradeCluster which made the WebSphere Application Server provisioning very efficient. We also pre-installed the EWLM Managed Server code on MSNAIX1, so provisioning the server into the EWLM management domain was also an efficient process. If we were to use Tivoli Provisioning Manager to provision the systems for us from the *bear mettle*, we would have to take all of these installation and configurations steps into considerations and create (or modify) workflows for each of these steps.
- The IBM Director eliminated the platform unique management approaches and provided us with real-time health monitoring of our servers and storage. Although we only defined a minimal amount of resources through the IBM Director, it allowed us to set known thresholds for our individual operating systems and monitor them through the Virtualization Engine Console. We were then able to relate these system thresholds, to our workload business goals. We found this to be a very powerful source for systems management.
- Because the EWLM Managed Servers send the performance data to the domain manager every 10 seconds and the EWLM Domain Manager aggregates those numbers every 30 seconds, there is a delay in the EWLM control center for the EWLM information. It took a little longer than we had expected to notice the effects of the newly provisioned system and the PI levels changing. One of these delays was caused by the interval preference we had set was too long (30 minutes) to determine a change we created in the short-term because the average of the 30 minute period was misleading. We shortened our interval to 5 minute (as well as 1 minute) to see the latest short-term effects of provisioning the second WebSphere Application Server server. Although with a mixed workload, this value might have been too short, for our testing purposes, it worked out just fine.
- We did not take into account that the default routing algorithm for the Web server was round-robin. Because the HTTP server still routed half of the transactions to the WebSphere Application Server which was running on a system which had other CPU intensive work, the transactions that were being routed there were still being effected by the CPU-consuming job. We could have, however, used Tivoli Provisioning Manager to provision a third WebSphere Application Server running on MSNAIX1 (or even a third AIX image) to offload even more of the transactions from MSNAIX2, we decided that our scenario produce the results that we were expecting, allowing us to use the EWLM control center, and the Virtualization Engine Console to monitor the workload and have Tivoli Provisioning Manager provision resources which were dynamically picked up by our management servers and middleware products, allowing us to achieve our goals.

12

How to fix hardware constraints using TPM and IBM Director

The following scenario depicts a situation where a spare xSeries BladeCenter HS20 Blade is provisioned with a Windows 2003 operating system image using Tivoli Provisioning Manager and IBM Director. This type of scenario can result from analyzing EWLM data and determining that an application requires additional hardware.

This chapter discusses the following topics:

- The topology for this scenario
- Configuring the IBM Director
- Configuring RDM
- Configuring Tivoli Provisioning Manager data center resources
- Configuring the Tivoli Provisioning Manager IBM-RDM.tcdriver and workflow
- Completing the scenario: Provisioning the spare Blade resource

Note: Remote Deployment Manager is not part of the Virtualization Engine. It is an add-on for IBM Director that facilitates the deployment of Windows and Linux operating systems to IBM Director managed servers. For more information, visit the following:

http://www-1.ibm.com/servers/eserver/xseries/systems_management/sys_migration/rdm.html

© Copyright IBM Corp. 2005. All rights reserved.

12.1 The topology for this scenario

Figure 12-1 describes the scenario logical environment.

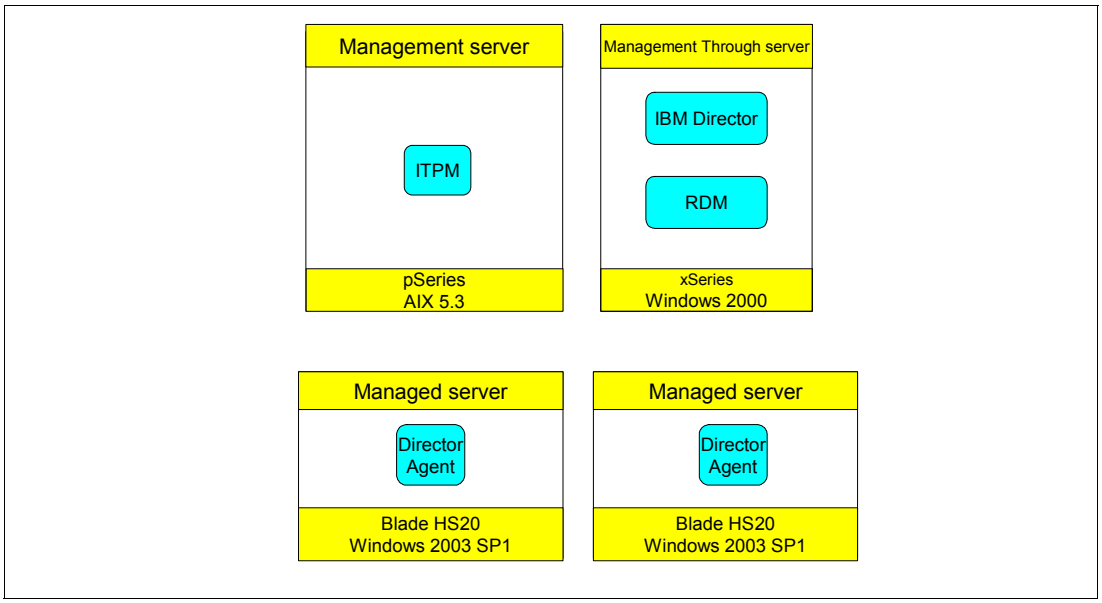

Figure 12-1 Scenario topology

The following lists describes the features of this topology:

- The management servers are:

  VEAIX43: 10.1.1.180
  Windows 2003: Tivoli Provisioning Manager 2.1

- The managed through servers are:

  VEWIN21: 10.1.1.178
  Windows 2000: IBM Director 4.20, Remote Deployment Manager 4.11 Update 3, MS DHCP Server (2 address scope for PXE boot. 10.1.1.200-201). RDM Windows Clone Install task named WIN2K3SP1

- The managed servers are:

  BLADEHS20_1 Provisioned with Windows 2003 image
  MSNWIN2 10.1.1.185 as the IP address

Note: Cygwin is a requirement for Tivoli Provisioning Manager running on Windows. It is also needed on Windows managed servers. Cygwin provides the proper command shell and tools that are used by Tivoli Provisioning Manager to communicate and manipulate its managed devices. In this scenario, the IBM Director system running on Windows has Cygwin installed and configured as described in the Tivoli Provisioning Manager 2.1.0 documentation.

The Tivoli Provisioning Manager DCM network information:

- Switch Fabric: ITSO_Switch - Ports 1-8
- VLAN: 1
- Subnetwork: ITSOVE - 10.1.1.0/255.255.255.0

12.2 Configuring the IBM Director

IBM Director Multiplaform was installed on our managed through using the Virtualization Engine installation wizard. For the purposes of this scenario, we added an additional component to IBM Director called Remote Deployment Manager (RDM)

Additionally, to enable Tivoli Provisioning Manager, IBM Director, and RDM functionality, we copies the following files from the Tivoli Provisioning Manager server to the IBM Director server:

- Copied %TIO_HOME%/bin/epprdmcli.jar to [Director Path]\classes
- Copied %TIO_HOME%/bin/RdmCli.TWGExt to [Director Path]\classes\extensions

After copying these files, we need to restart the IBM Director Server's service for these additions to load using the following commands:

```
net stop twgserver
net start twgserver
```

The following sections detail the configuration of Remote Deployment Manager. As mentioned earlier, RDM is not part of the Virtualization Engine suite of products and many of its functions are outside of the scope of this book. We focus mainly on the functions that we need to facilitate the operating system provisioning using Tivoli Provisioning Manager.

12.3 Configuring RDM

This section describes the configuration and the Windows Clone Install task. This task performs the actual operating system installation on our spare blade server when invoked by the Tivoli Provisioning Manager workflow that are described later in this chapter.

Note: The Windows 2003 Server operating system image we are deploying to our spare server was built using the Get Donor RDM task in IBM Director. The process for installing Windows 2003 Server on the donor system, preparing the installation for duplication using the Microsoft System Preparation Tool (Sysprep), and obtaining the image using Get Donor are outside the scope of this discussion. Refer to the RDM Operations Guide for more information about preparing a system image. This documentation is available at:

http://ats.greenock.uk.ibm.com/Blade/RDM/rdm411compatibilityguide.pdf

Configuring the Windows Clone Task

To configure the Windows Clone Task, do the following:

1. From the IBM Director Console, expand the RDM container on the task pane
2. Right-click the Windows Clone Install container and select **Create New Task**.
3. Specify a name for the new task. Remember this RDM task name because you will use it later during the Tivoli Provisioning Manager configuration.
4. Click **Next**.
5. Specify the location of the server image.
6. Click **Next**.
7. Specify the Name, Organization, and Product Key that are associated with the Windows installation.
8. Click **Next**.

9. Choose the Licensing Type. For our purposes, we use Per Server licensing mode.
10. Click **Next**.
11. Choose Regional Settings, and click **Next**.
12. Then specify the workgroup or domain name. We choose a workgroup named WORKGROUP.
13. Click **Next**, then specify the network protocols used. TCP/IP is selected by default.
14. Click **Next** again, and specify the TCP/IP configuration. We configure an IP address manually that matches how our BLADEHS20_1 server is listed in the Tivoli Provisioning Manager configuration that we will do later.
15. Click **Next** and the task is configured.

You can edit the task at any time by right-clicking it and choosing **Edit Task**.

At this point the RDM Windows Clone Install task is prepared for deployment to any of our IBM Director managed systems. The next section discusses the configuration that is required to prepare Tivoli Provisioning Manager to run RDM Server workflow for provisioning servers through IBM Director.

12.4 Configuring Tivoli Provisioning Manager data center resources

Our management server is running AIX 5.3 (64-bit) and Tivoli Provisioning Manager 2.1.0.2. One of the core components of Tivoli Provisioning Manager is the data center model (DCM). The DCM is the mapping of the IT infrastructure, including objects such as boot servers, server pools, clusters, network devices, and software information. The DCM provides the overall view of what the IT infrastructure “looks like”.

The DCM is populated manually using the Tivoli Provisioning Manager Web interface. Alternatively the DCM can be imported as an XML file, or by pulling information out of a Tivoli NetView map using the Discovery.tcdriver and associated NetView workflow.

For the purposes of this scenario we are using a combined approach. The network switch, VLAN, subnetwork and Tivoli Provisioning Manager server entries were imported into the DCM using an XML file as in Example 12-1. The remaining information, including Resource Pools, Servers, and Software Stacks, was added manually using the Tivoli Provisioning Manager Web interface.

Example 12-1 DCM XML configuration

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- edited with XMLSPY v5 rel. 4 U (http://www.xmlspy.com) by IBM Corporation (IBM
Corporation) -->
<!DOCTYPE datacenter SYSTEM "xmlimport.dtd">
<datacenter>
   <switch-fabric name="ITSO_Fabric" />
   <switch name="ITSO_Switch" locale="en_US" fabric="ITSO_Fabric" failed="false"
in-maintenance="false">
      <switch-module>
         <switch-port failed="false" managed="false" netboot-enabled="false" vlan="1"
port-number="1" enabled="true" layer1-interface-type="unknown" />
         <switch-port failed="false" managed="false" netboot-enabled="false" vlan="1"
port-number="2" enabled="true" layer1-interface-type="unknown" />
```

```
         <switch-port failed="false" managed="false" netboot-enabled="false" vlan="1"
port-number="3" enabled="true" layer1-interface-type="unknown" />
         <switch-port failed="false" managed="false" netboot-enabled="false" vlan="1"
port-number="4" enabled="true" layer1-interface-type="unknown" />
         <switch-port failed="false" managed="false" netboot-enabled="false" vlan="1"
port-number="5" enabled="true" layer1-interface-type="unknown" />
         <switch-port failed="false" managed="false" netboot-enabled="false" vlan="1"
port-number="6" enabled="true" layer1-interface-type="unknown" />
         <switch-port failed="false" managed="false" netboot-enabled="false" vlan="1"
port-number="7" enabled="true" layer1-interface-type="unknown" />
         <switch-port failed="false" managed="false" netboot-enabled="false" vlan="1"
port-number="8" enabled="true" layer1-interface-type="unknown" />
      </switch-module>
      <switch-vlan vlan-number="1" />
   </switch>
   <subnetwork name="itso_ve" ipaddress="10.1.1.0" netmask="255.255.255.0">
      <vlan vlan-number="1" fabric="ITSO_Fabric" />
   </subnetwork>
<spare-pool name="TPM Server" locale="en_US" os-type=" AIX" vlan="1" fabric="ITSO_Fabric">
      <server name="veaix42" locale="en_US">
         <nic name="en0" macaddress="7241E0003002" netboot-enabled="false" managed="true"
connected-to-switch="ITSO_Switch" connected-to-port="1">
            <network-interface name="Local Area Connection" ipaddress="10.1.1.187"
netmask="255.255.255.0" />
         </nic>
```

You can find information about importing a DCM in 6.5.8, “Steps 7 to 9: Post-installation tasks” on page 158.

12.4.1 Creating a resource pool

The Tivoli Provisioning Manager Web interface contains two frames.

- The left-hand frame represents the navigation bar which contains the following tabs:
  - Data center assets and resources
  - System configuration and workflow management
  - Customer applications
- The right-hand frame provides the focused view of the items that are selected from the navigation bar as shown in Figure 12-2.

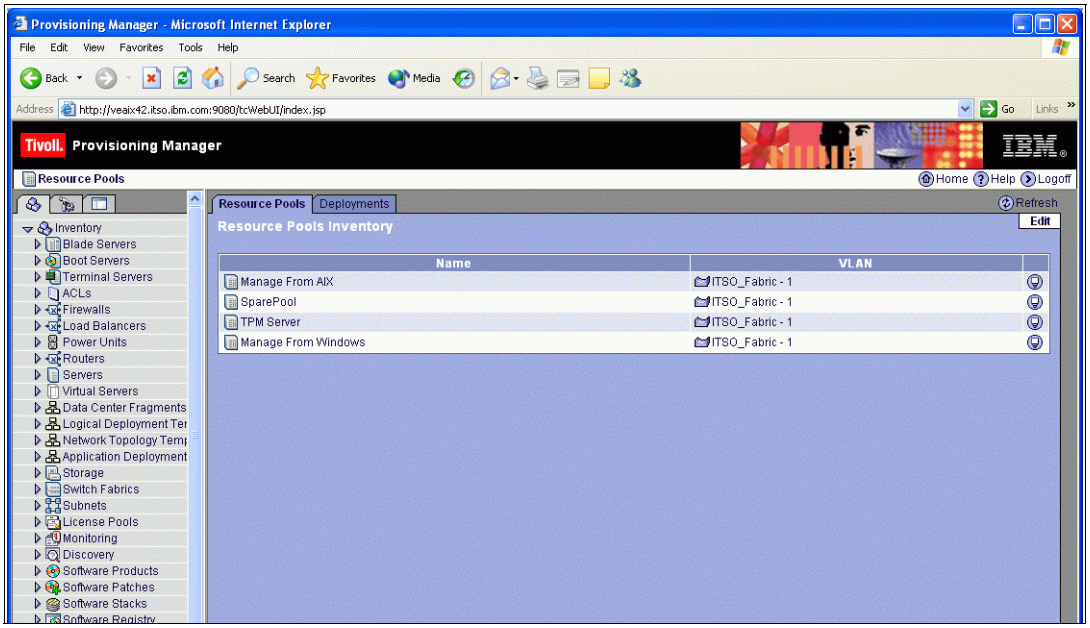

Figure 12-2 Tivoli Provisioning Manager *interface showing data center Resource Pools*

To add additional resource pools, do the following:

1. Click **Edit** and then **Add Resource Pool**, as shown in Figure 12-3.

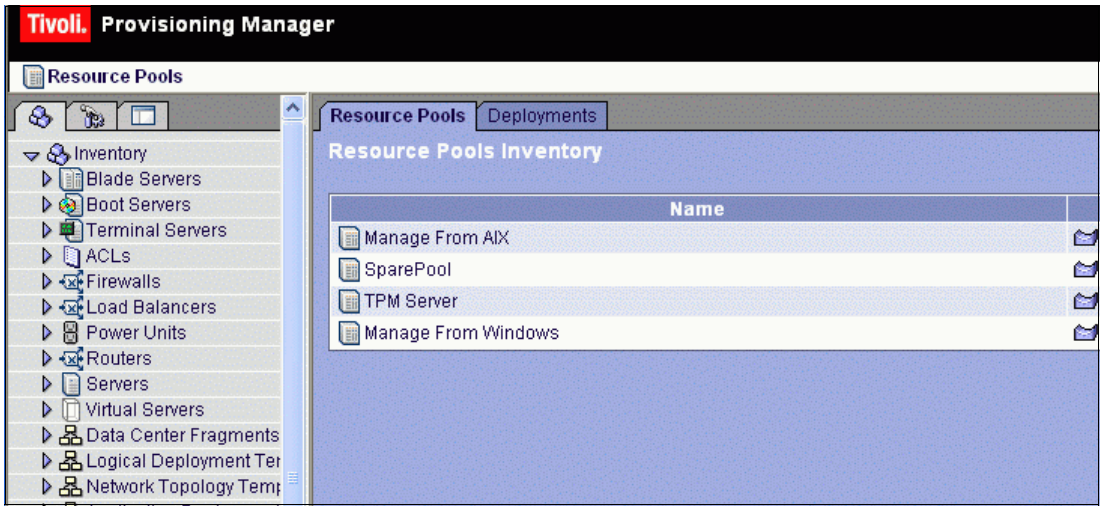

Figure 12-3 Adding a resource pool

2. Specify the name of the Resource Pool, the VLAN in which the resources in this pool are located, as well as the locale for the environment, as shown in Figure 12-4.

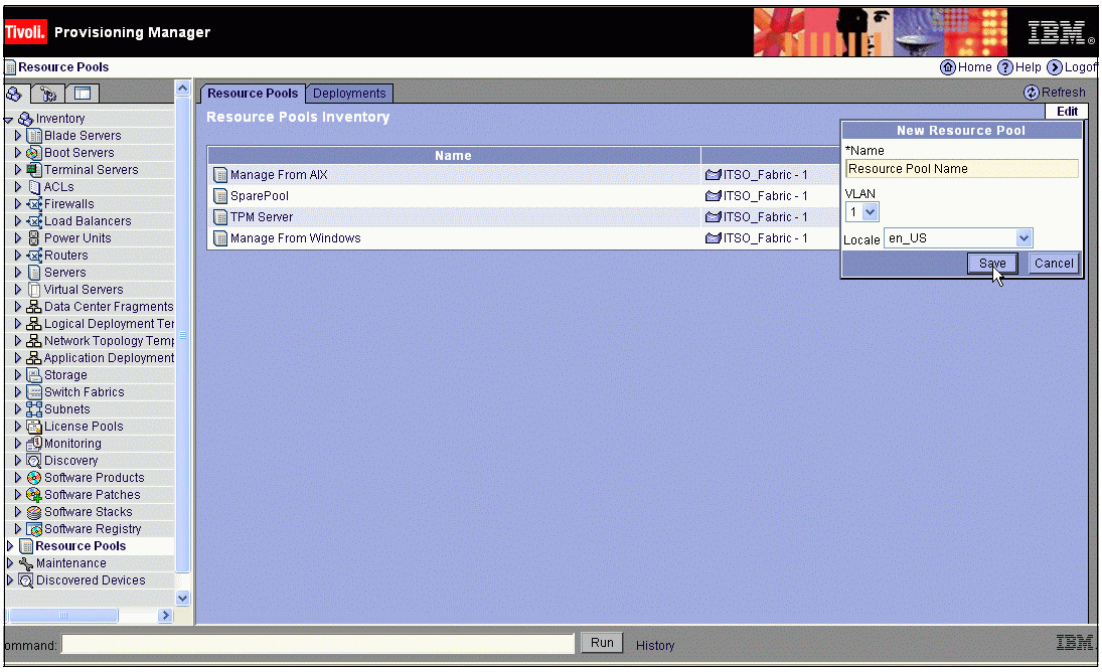

Figure 12-4 Configuring the resource pool

The two resource pools that we create in our scenario are:

- Spare Pool, which represents our unallocated server resources.
- Managed From Windows, which represents our production Windows server environment

12.4.2 Creating the boot server

The boot server in this scenario is the server that is running IBM Director and RDM, also known as the *managed through* server. The approach to adding a boot server is similar to adding a resource pool. Do the following:

1. From the navigation bar on the left frame of the Tivoli Provisioning Manager Web interface choose the Data center assets and resources tab and click **Boot Servers**.
2. On the right frame, click **Edit**, and then **Add Boot Server**. Then specify the name and locale for the boot server and click **Save**.
3. Select the boot server that you added.
4. Configure a NIC by clicking **Edit**, and then **Add NIC**. Complete the appropriate fields as shown in Figure 5 on page 298.

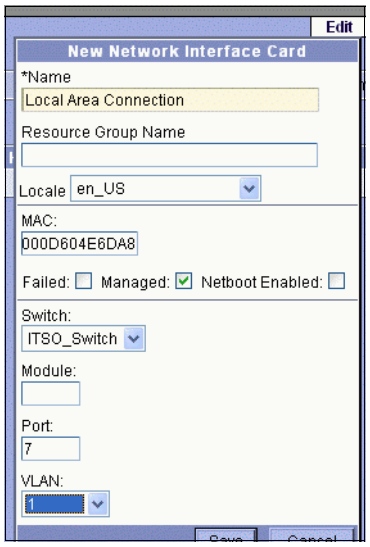

Figure 12-5 Adding a NIC in Tivoli Provisioning Manager

Tip: If you specify a port that is currently allocated, an error message will appear:

```
COPJEE100E The system cannot connect the NIC to port 7 because the port is being used
by another device.
COPJEE272E (Problem ID: UI548363).
```

You might need to navigate to your switch to determine an available port number. This information is located by clicking “Switch Fabrics” from the navigator bar and clicking the appropriate switch.

5. Click **Save** to close the window.
6. Click the circular icon next to the newly added NIC, and choose **Add Interface** as shown in Figure 12-6.

Figure 12-6 Adding an interface

7. Complete the appropriate fields, as shown in Figure 12-7 , then click **Save**.

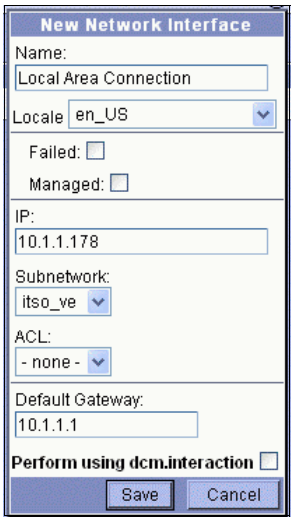

Figure 12-7 The new interface properties

12.4.3 Creating the spare Blade

Our spare server is created within the Servers container on the navigator bar. By clicking the arrow next to Servers, we see all the resource pools that are defined in our DCM. For example, the two resource pools that we created earlier, Spare Pool and Managed From Windows, are listed, as shown in Figure 12-8. In our scenario, the spare server is added to the Spare Pool container.

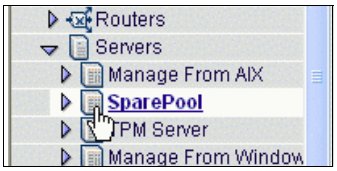

Figure 12-8 Creating the spare Blade

To add the spare server resource, click **Spare Pool**. Then, add the spare server resource using the same procedures that you used to create the boot server.

- Add the server
- Add a NIC to the server
- Add an Interface to the NIC

12.4.4 Creating the software stack

The software stack is the Tivoli Provisioning Manager representation of our Windows 2003 server image. This image exists within our RDM configuration in IBM Director as described earlier. The Tivoli Provisioning Manager workflow that we use communicate with RDM and invoke the RDM Windows Clone Install task are described in 12.6, “Completing the scenario: Provisioning the spare Blade resource” on page 306.

The following steps are used to create the Software Stack data center object in Tivoli Provisioning Manager:

1. From the navigation bar on the left frame of the Tivoli Provisioning Manager Web interface choose the Data center assets and resources tab and click **Software Stacks**.

2. From the right-hand pane, choose **Edit**, and then **Add Software Stack**.
3. Provide a name for the software stack. For our purposes we use WIN2K3SP1, which matches our Windows image as it appears in the RDM Windows Clone Install.
4. Click **Save** and the new software stack is listed.
5. Select the newly added software stack to make configuration changes.
6. The top section of the General tab contains two drop down menus. Modify these as follows:
   - Software Stack Type: Image Software Stack
   - Device Driver: Image Software Stack
7. Click **Change**.
8. Continuing on the General tab, under the Image Stack Configuration section, change the Boot Server to the boot server that you created earlier and click **Save**.

Figure 12-9 shows our example.

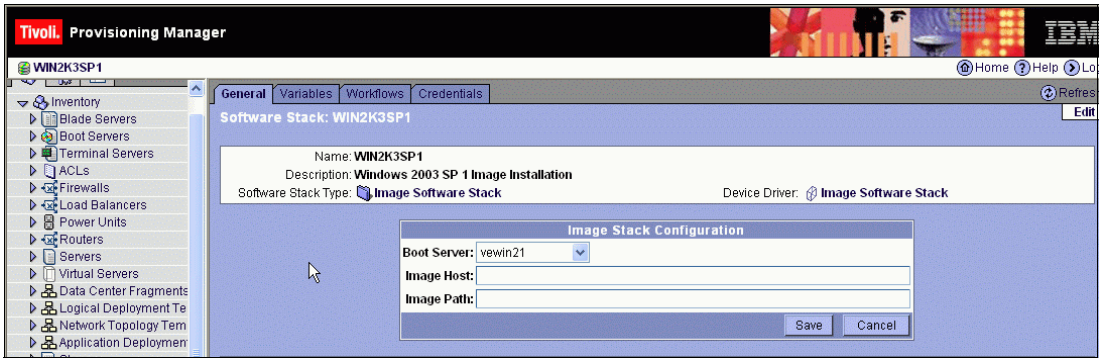

Figure 12-9 Creating the software stack

The completes the core data center resource configuration for the devices that are involved in our scenario. The next section discusses the configuration of the IBM-RDM.tcdriver. We revisit these data center resources for additional configuration.

12.5 Configuring the Tivoli Provisioning Manager IBM-RDM.tcdriver and workflow

The automation power of Tivoli Provisioning Manager lies in is its ability to execute a wide range of scripts which it is called workflow. These workflow range from moving resources in between resource pools, to configuring network devices, to installing software on servers. Tivoli Provisioning Manager comes pre-packaged with several workflow. These workflow can be fully edited and customized, or workflow can be created from the ground up using Tivoli Provisioning Manager's scripting language and java plug-in's.

Another automation component provided by Tivoli Provisioning Manager is the device driver. A device driver is essentially a manufacturer provided collection of workflow, scripts, installation instructions that can be used to drive a specific process.

In our scenario, we work with an IBM device driver that contains workflow and shell scripts which allow Tivoli Provisioning Manager to communicate with IBM Director and Remote Deployment Manager. Through this driver Tivoli Provisioning Manager can initiate an RDM

deployment task that provisions our spare server resource with a new Windows 2003 operating system image.

The Tivoli Provisioning Manager driver we are implementing is called “RDM Server” and is packaged as the “IBM Automation Package for Remote Deployment Manager (RDM)” or IBM-RDM.tcdriver.

You can download the latest Tivoli Provisioning Manager version specific IBM-RDM.tcdriver from the On Demand Automation Catalog Web site at:

http://www-18.lotus.com/wps/portal/automation

In our scenario, we are using IBM-RDM.tcdriver version 2.3.2. This driver comes with the version of Tivoli Provisioning Manager (version 2.1.0 Fix Pack 2) that we are using. If you have the base version of Tivoli Provisioning Manager 2.1.0, you can download this file and install it on your system. The next section describes the process of installing a new driver in the event that this applies.

12.5.1 Installing or updating the IBM-RDM.tcdriver

Tivoli Provisioning Manager provides a shell script for installing device drivers called **tc-driver-manager.sh**. Example 12-2 contains the script options. The default location for the Tivoli Provisioning Manager scripts is:

- AIX and Linux: install_dir/IBM/tivoli/thinkcontrol/tools
- Windows: install_dir\IBM\tivoli\thinkcontrol\tools

Example 12-2 tc-driver-manager.sh command

```
-bash-2.05b$ tc-driver-manager.sh
TCDrivermanager was started.
Config directory = 'file:/opt/IBM/tivoli/thinkcontrol/config/'.
Driver directory = '/opt/IBM/tivoli/thinkcontrol/drivers/'.
java.lang.ArrayIndexOutOfBoundsException: 0

Available TCDriverManager commands are:

   forceInstallDriver <driverName> [options]
   forceUnInstallDriver <driverName> [options]
   getDescription <driverName> [options]
   getDocumentation <driverName> [options]
   getDriverStatus <driverName> [options]
   installDriver <driverName> [options]
   installNoItems <driverName> [options]
   listAllStr [options]
   listDeviceModels [options]
   listInstalledDeviceModels [options]
   uninstallAllDrivers
   uninstallDriver <driverName> [options]
   options:
      -fod         fails driver install whenever
                   there are references to the driver <default=true>
      -overwrite   if overwrite of files allowed during driver
                   installation <default=false>
```

To install the IBM-RDM.tcdriver, run the following command:

```
tc-driver-manager.sh installDriver IBM-RDM.tcdriver
```

The driver is unpackaged into the appropriate directories making it available for use through the Tivoli Provisioning Manager Web console.

Note: For an Tivoli Provisioning Manager server running on Windows, either run this command within the Cygwin environment or use the Windows command script, **`tc-driver-manager.cmd`**, which is also included as part of the base Tivoli Provisioning Manager installation.

With the driver installed, we can now configure the data center resources to use its resource.

12.5.2 Configuring workflow and variables on the data center resources

In this section, we revisit three of these resources for additional configuration.

- Boot server
- Spare Blade
- Software stack

We also make some configuration changes to two other resources:

- Tivoli Provisioning Manager Server - VEAIX42
- The Spare Pool resource pool

Boot server

To configure the boot server:

1. From the navigation bar, expand the Boot Servers container and click the boot server that you added earlier.
2. Go to the Variables tab that is associated with the boot server.
3. Add the variable by entering information in the appropriate field, as shown in Table 12-1, and then click **Add**.

Table 12-1 Boot server variables

Key	Component	Value
IBMDirectorScriptsPath	Entire System	C:/cygwin/home/ibmdirector/IBMDirectorScripts/

Note: The Value in Table 12-1 corresponds to the IBMDirectorScripts directory, which is located on the IBM Director Server and must be created manually.

4. Go to the Workflow tab and set the device driver to RDM Server, and click **Change Driver**.
5. Go to the Credentials tab and configure Service Access Points and Credentials for the boot server.
6. Click **Edit** → **Add Access Point** and do the following:
   - In the Name field, enter `IBMDirector`.
   - In the Type field, select **other / LOCAL-EXEC**.
   - Select **Authentication**.
   - Ensure that Host is *not* selected.
   - Set the locale to `en_US`.
7. Click **Save**.

8. Click the IBMDirector Service Access Point edit icon and select **Add Credentials: password**. Do the following:
   - In the Search Key field, enter `IBMDirectorCLI` .
   - In the User Name field, enter the ID that has administrative access to IBM Director on the IBM Director Management Server.
   - In the Passphrase field, enter the password for the entered user ID.
   - Leave the Enable Password and Key Fingerprint fields blank.
9. Click **Save**.
10. Click **Edit**, and then **Add Access Point**.
    - In the Name field, enter: `SSH`.
    - In the Type field, select **ipv4 / SSH**.
    - Select **Authentication**.
    - Ensure that Host is selected.
    - Set the locale to `en_US`.
11. Click **Save**.
12. Click the SSH Service Access Point edit icon and select **Add Credentials: RSA**.
    - In the Search Key field, enter `executecommand`.
    - In the User Name field, enter the ID that has administrative access to IBM Director on the IBM Director Management Server.
13. Click **Save**.
14. Click that icon to the right of the credential and set it to default.
15. Click **SSH Service Access Point** on the Workflow tab, then click **Logical Device Operation**. Select **-all workflows**. Attach the following workflows:
    - Default_Device_Execute_Command
    - SCP_Copy_File
    - SSH_Ping_IP_Address.
    - SSH_RSA_Execute_Command
16. Go to the Credentials tab, and assign the SSH Service Access Point to the Device Operations execute-command and file-transfer.

Spare Blade

To configure the spare Blade, do the following:

1. From the navigation bar, expand Servers, then Spare Pool, and click the spare server resource, BLADEHS20_01.
2. Go to the Variable tab and add the variable as shown in Table 12-2.

Table 12-2 Adding a variable into the spare pool

Key	Component	Value
UUID	Entire System	422EA9DADF1DB211A72B5CDA030C8B48

The system UUID is obtained from IBM Director as shown in Figure 12-10. For non-IBM systems use the MAC address as the UUID.

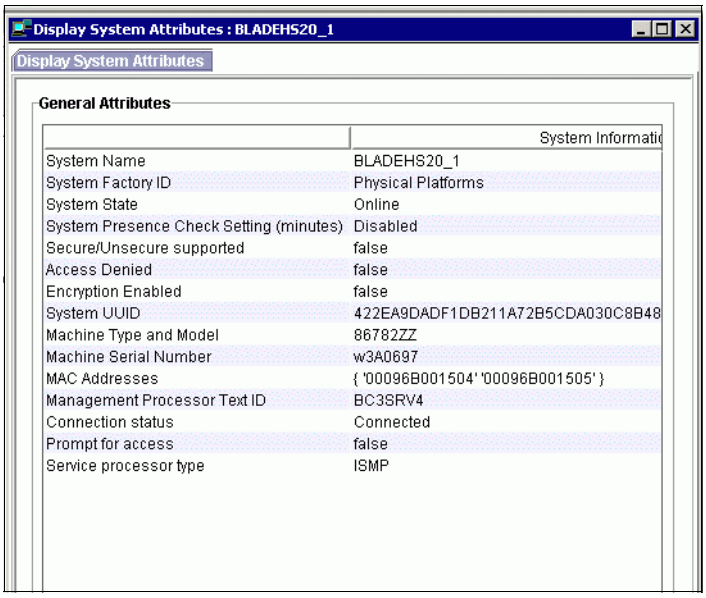

Figure 12-10 System attributes in IBM Director

Software stack

To configure the software stack, do the following:

1. From the navigation bar, expand the Software Stack container and click the software stack that you added earlier.
2. Go to the Variables tab and add the variables listed in Table 12-3.

Table 12-3 WIN2K3SP1 Software Stack variables

Key	Component	Value
RDMTaskName	Entire System	WIN2K3SP1
TimeOut	Entire System	240
SetIP	Entire System	True

Note: The value for RDMTaskName *must* be *identical* to the Windows Clone Install task in RDM.

The Tivoli Provisioning Manager server

Our Tivoli Provisioning Manager server is listed under the Servers container in the Tivoli Provisioning Manager Server resource pool. If you have not already added a Tivoli Provisioning Manager server, do this using the procedures for creating the boot server as described in “Boot server” on page 302.

To configure the Tivoli Provisioning Manager server, do the following:

1. From the navigaton bar, expand Servers and Tivoli Provisioning Manager Server. Then, select the Tivoli Provisioning Manager server, VEAIX42.
2. Go to the Credentials tab, and add a Service Access Point.

3. Click **Edit** → **Add Access Point** and do the following:
   - In the Name field, enter `LocalExc`.
   - In the Type field, select **other / LOCAL-EXEC**.
   - Select **Authentication**.
   - Ensure that Host is selected.
4. Click **Save**.

   Note: This Service Access Point does not require credentials

5. Add another Service Access Point by doing the following:
   - Click **Edit** → **Add Access Point**.
   - In the Name field, enter `SSH`.
   - In the Type field, select **ipv4 / SSH**.
   - Select **Authentication**.
   - Ensure that Host is *not* selected.
6. Click **Save**. Then, do the following:
   - Click the Service Access Point edit icon and select **Add Credentials: RSA**.
   - In the Search Key field, enter `executecommand`.
   - In the User Name field, enter the ID that has administrative access to the Tivoli Provisioning Manager server.
7. Click **Save**.
8. Click the icon to the right of this credential and set it to default.
9. Click the created LocalExc Service Access Point and on Workflow tab and attach the device driver Deployment Engine Service Access Point.
10. Go to the Credentials tab, assign the LocalExc Service Access Point to Device Operation execute-command, and do the following:
    - Click the created SSH Service Access Point, and on Workflow tab, click **Logical Device Operation**. Select **-all workflow**.
    - Select the workflow:
      - Default_Device_Execute_Command.
      - SCP_Copy_File
      - SSH_Ping_IP_Address.
      - SSH_RSA_Execute_Command

Configure the resource pool

To configure the resource pool, do the following:

1. From the navigation bar, click **Servers** and then **Managed From Windows**. This is our production server pool.
2. Go to the Workflow tab and click **Logical Device Operation**. Select **-all workflow** and select the workflow Default_Pool_Initialize_Server.
3. Go to the Software Stack tab, click the drop-down menu, and choose the WIN2K3SP1 software stack.
4. Click **Save**.

This concludes the IBM-RDM.tcdriver configuration. The environment is now prepared for server provisioning. In the next section, we describe the provisioning process.

12.6 Completing the scenario: Provisioning the spare Blade resource

As described in the beginning of this chapter, our scenario depicts a situation where a spare xSeries BladeCenter HS20 Blade is provisioned with a Windows 2003 operating system image using Tivoli Provisioning Manager and IBM Director.

We have categorized this as a situation where the environment has experienced some hardware constraints and requires additional hardware resources in order to meet business goals. This type of scenario can result from analyzing EWLM data and determining that an application requires additional hardware. Or, by identifying IBM Director alerts though the Virtualization Engine console related to disk or memory constraints. This type of scenario also applies to instances where physical hardware issues occur and new servers are required to quickly replace failing or failed systems.

In an overall sense, it is important to keep in mind that our scenario is something of an oversimplification of a provisioning scenario. This scenario should provide a taste of the Tivoli Provisioning Manager and IBM Director capabilities.

The core function implemented by the IBM-RDM.tcdriver in our provisioning scenario is the RDM_Server_Image_Install workflow. This workflow is assigned to our boot server and initiates the interaction with our IBM Director and RDM server.

The following sections lists the workflow and scripts that are involved in implementing the IBM-RDM.tcdriver.

Workflows

Workflows are located in the %tio_home%/workflow/ directory and are as follows:

- RDM_Server_Image_Install.wkf - Workflow that installs an image using RDM, first updating the IP address.
- IBM_Director_Run_Task.wkf - Uses IBM Director to run a non-interactive task.
- IBM_Director_Run_Task_Helper.wkf - Executes local command.
- IBM_Director_Wait_For_Completion.wkf - Polls IBM Director until a task has completed.
- IBM_Director_Wait_For_Completion_Helper.wkf - Executes local command.
- IBM_Director_RDM_Update_IP.wkf - Updates IP address for an install image.
- IBM_Director_RDM_Update_IP_Helper.wkf - Executes local command.
- IBM_Director_Transfer_DirCmd2_scripts.wkf - Copies script files from ITITO server to Boot Server.
- IBM_Director_Copy_files_SSH-SCP.wkf - Copies files via SCP.
- NLSString.wkf - Finds locale of a Boot Server and returns the correct translation of Power Off Now IBM Director non-interactive task.

Scripts

Scripts are located in the %tio_home%/bin/ directory and include the following:

- DirRunTask_DC2.sh - Script to run a non-interactive task on the IBM Director Management Server
- DirTaskStatus_DC2.sh - Script to get status on a task running on IBM Director Management Server

- DirRDMSetIP_DC2.sh - Script to run an RDM database update on the IP field of the install task

To provision our spare Blade server, do the following:

1. From the navigator bar expand the Servers container, then click the Spare Pool resource group.
2. Our spare blade, BLADEHS20_1, is listed in the list of spare pools. It is currently offline, marked in `Maintenance` mode as shown in Figure 12-11.

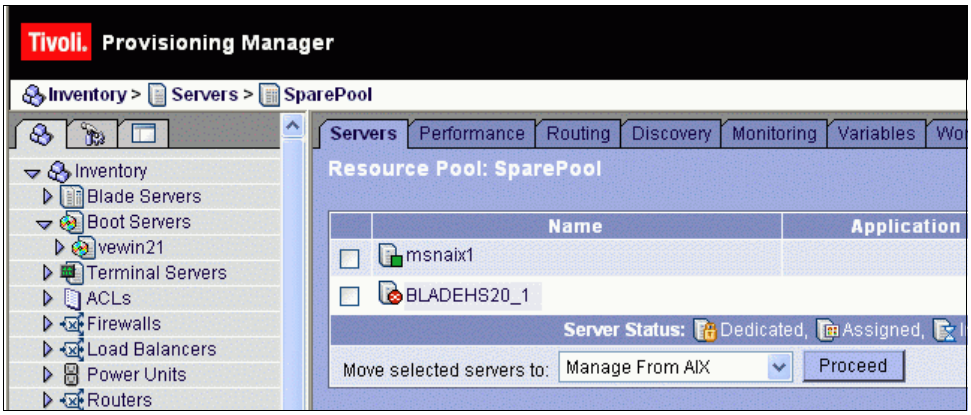

Figure 12-11 Spare pool servers list

3. Select **BLADEHS20_1**.
4. Expand the drop-down menu next to **Move selected servers to:** and specify Managed From Windows. This represents our production Windows server pool.
5. This immediately initiates the Default_Pool_Intialize workflow assigned to the Managed From Windows resource pool.
6. During this initialization, the WIN2K3SP1 software stack, also assigned to the Managed From Windows resource pool, is applied.
7. The application of this software stack then kicks off the RDM_Server_Image_Install assigned to our boot server and the process of provisioning our server begins.
8. At this point the BLADESH20_1 server is listed in the Managed From Windows resource pool in a transitioning state while the RDM_Server_Image_Install workflow executes.

The RDM_Server_Image_Install workflow executes as follows:

1. SCP_Copy_File copies the three .sh scripts to the IBMDirectorScripts folder on the director server.
2. SSH_RSA_Execute_Command executes the .sh scripts that were copied to the IBMDirectorScripts folder on the director server.
3. If you look at the Director Console, you see that a task has been initiated against the BLADEHS20_1 server. This indicates that the WIN2K3SP1 Windows Clone Task has started.
4. The server first waits for a wake on LAN call to perform a network (PXE) boot.
5. When this occurs the image installation process begins.
6. When the image is delivered successfully, the task completes.

7. The IBM_Director_Wait_For_Completion workflow reports the task completion to Tivoli Provisioning Manager, and the BLADESH20_1 server in the Managed From Windows resource pool is ready.

At this point the system has been provisioned.

13

How to optimize an xSeries IT environment using IBM Director and Virtual Machine Manager

This chapter discusses the collaboration between the IBM Director and a VMware environment that is composed of a VMware VirtualCenter and a VMware ESX Server. The vehicle to make this collaboration become reality is the IBM Virtual Machine Manager (VMM). We used VMM Version 1.0 for this book.

Documentation is available all the individual components. The intention of this chapter is not to replicate that document. It is more to describe the value when combining the products together. Therefore, we knowingly use some replication which is necessary to get a common understanding of the products and which is needed to explain the workflow in our scenario.

For other cases, if you need a more detailed explanation of the product's functional range and how to use the samples, refer to the product documentation.

© Copyright IBM Corp. 2005. All rights reserved.

13.1 IBM Director enhancement through VMM

This chapter uses the installation of the IBM Director on Windows (see Chapter 7, “Installing the Virtualization Engine console, IBM Director, TPM, and EWLM on Blade with Windows” on page 161) and the VMM installation (see 9.3, “The VMM installation process” on page 229). This section describes how VMM enhances IBM Director in using the virtualization software. VMware in combination of VMware VirtualCenter and VMware ESX Server.

Note: With VMM Version 1.0, you can use Microsoft Virtual Server to build a management environment. However, in our scenario, we focus on the VMware implementation.

From an IBM Director perspective, VMM adds new objects and functions to the base installation. Figure 13-1 shows a visual representation of the object structure.

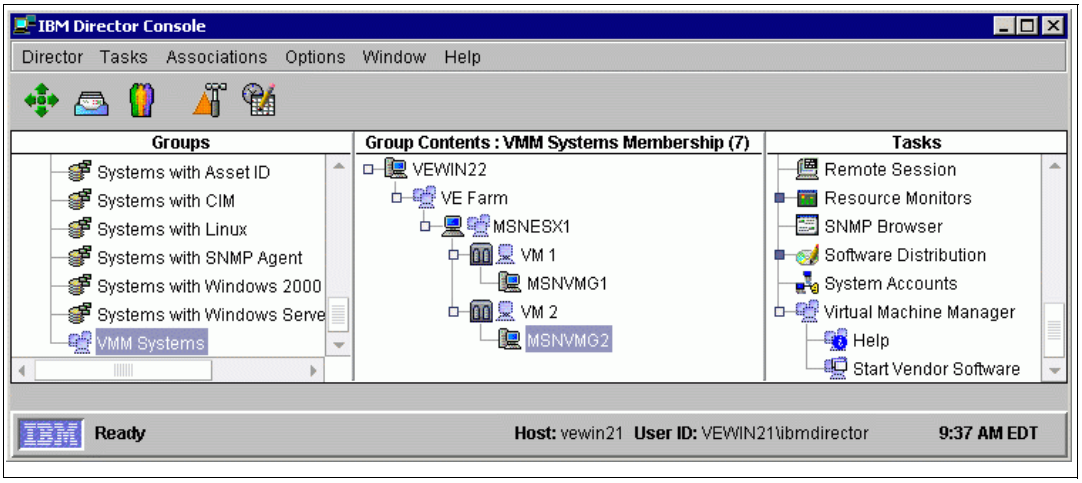

Figure 13-1 Selected group VMM Systems and corresponding Group Contents view of VMM Systems

The following lists additional objects and functions that are used within the examples in this chapter:

- The coordinator object in our environment is VEWIN22.

  A *coordinator object* is a node in the management domain that has VMware VirtualCenter server plus VirtualCenter Web Service, IBM Director Agent and VMM installed and running. Therefore, there could be a range of VMware VirtualCenter in the environment, but we use only one in our environment.

  All objects that are managed trough a VMware VirtualCenter are bequeath to IBM Director because of the linkage between VMM and VMware. Consequently we see all farms, VMware ESX Servers, virtual machines and guest operating systems that the VMware VirtualCenter manages.

  In our environment, we also have IBM Director control on VEWIN22 because of the installed IBM Director Agent.

  The functions on a coordinator object are:

  – Enter Credentials
  – Revoke Credentials
  – Discover Farms
  – Lookup VMM Coordinator Object Attributes
  – And standard IBM Director Power Management™ functions

- The farm in our environment is Virtualization Engine Farm.

  A *farm* is a logical organizational form of VMware VirtualCenter to build a logical structure of the management domain. A way to use the farm construct is to group hosts and define unique names such as development, test, and production. Virtualization Engine Farm stands for Virtualization Engine Farm.

  The management capabilities of farms are based on logical operations such as add or remove a farm and add or remove a host to a farm.

  The IBM Director Console VMM direct management functions on a farm object are:

  - Add Host to a Farm
  - Lookup VMM Farm object attributes
  - There is no IBM Director Power Management functions available because a farm object is only an organizational object

- The host in our environment is MSNESX1.

  A *host* could be either a VMware ESX Server or a Microsoft Virtual Server. In our environment we have one VMware ESX Server implemented with the server name MSNESX1 on 9.12.4.192.

  On MSNESX1 the IBM Director Agent for Linux is installed and we manage this server as a regular Linux entity in the IBM Director domain.

  The functions on a host object are:

  - Remove Host From Farm
  - Discover Virtual Machines
  - Power On All Stopped Virtual Machines
  - Force Power Off All Running Virtual Machines
  - Suspend All Running Virtual Machines
  - Resume All Suspended Virtual Machines
  - Lookup VMM Host Object Attributes
  - And standard IBM Director Power Management functions

- The virtual machines in our environment are VM 1 and VM 2.

  A *virtual machine* is a logical abstraction of a physical hardware to make the virtual machine independent of unique hardware requirements with the idea of an easy move across heterogeneous hardware platforms.

  With the virtual machine we have a physical hardware container that we could ‘fill’ with an operating system. A virtual machine for itself has no function therefore the only management capability resides on the power management that ideally reflects the hardware container of a server.

  The functions on a virtual machine object are:

  - Lookup VMM Host Object Attributes
  - And standard IBM Director Power Management functions

- The guest operating systems in our environment are MSNVMG1 and MSNVMG2.

  The *guest operating system* is a software container from the IBM Director perspective that contains everything needed for a server, including the IBM Director Agent.

In summary, we can say that the usage of the IBM Director agents and the VMM enhancement creates a holistic representation of all objects in the IBM Director managing domain.

13.2 The topology

The linchpin in the usage of IBM Director in our environment is to offer an overview of which are native IBM Director functions and which one are only available through the implementation of VMM in that IBM Director and VMware ESX Server play. The following sections show some of those native functions and those that the VMM enhancement makes available.

IBM Director native functions

IBM Director does not differentiate between a physical server (such as VEWIN21 in our environment) or a virtual server (such as MSNVMG1) if the IBM Director Agent is installed on a virtual server. They behave the same and provides the same functions, such as:

- Power Management: Shutdown and power off (now), restart (now), suspend.
- Remote Session: Start a session at the corresponding system, such as a Windows command prompt on Windows systems.
- Resource Monitor: Get all kind of system monitors, processor utilization, and definition of thresholds on a system.
- System Accounts: Manage the system accounts on a give server without the need to launch the corresponding system account application on that server. IBM Director is providing the interface.

> **Note:** The VMM object Farm is a logical construct to map the VMware VirtualCenter organizational structure equivalent to IBM Director. The Farm object represents no IBM Director manageable object and therefore power management and other tasks and operations could not apply to.

IBM Director VMM enhanced functions

On one hand, there is a function set that is available with whom the IBM Director Console user could launch VMware VirtualCenter specific tasks and on the other hand you get one function on VMM objects to lookup VMM Object attributes. Depending on the object and his function, different attributes are supported. The following list provides an overview of what functions are supported and which VMM object attributes could be retrieved:

- Coordinator Management
  - Enter or revoke credentials for a given VMware VirtualCenter management domain
  - Discover Farms
  - VMM objects: Library Vendor, Library Version, Vendor URL, TCP/IP Addresses, TCP/IP Hosts
- Farm Management
  - Add a host to IBM Director
  - VMM objects:Host Count, Virtual Machine Count, ID
- Host Management
  - Add or remove a host to IBM Director
  - Discover virtual machines
  - Power on/off/suspend/resume functions on one defined virtual machine
  - Power on/off/suspend/resume functions on all virtual machines running on a VMware ESX Server
  - VMM objects: System UUID, TCP/IP Hosts, Type of Hypervisor, Virtual Machine Count, Physical CPU, ID

- Virtual machine management

  VMM objects: Virtual CPU, Vendor Display Name, System UUID, Disk undoable behavior

13.3 Some scenarios that make use of the VMM enhancements

There are two different ways to make use of the VMM enhancements:

- The first way is the *manual* way to use the IBM Director console user interface to run commands interactively.
- Another way is the *automatic* way to use the server function of IBM Director. For instance to build actions and expanding actions could be launched through a scheduler with many variations on constrains and define thresholds to trigger events that launches actions and so forth.

Another goal is to show that there is no boundary between a real server and a virtual server. The current version of VMM 1.0 has some limitations that do not allow a direct manipulation on some VMware VirtualCenter or Microsoft Virtual Server objects such as creating a new virtual machine. If you plan to use those functions you have to use the native management interface from the vendors to deal with it.

13.3.1 Topology of the physical and logical servers

To better understand the scenarios and their context, we introduce the different physical and logical servers and the terminology that we use in that context (the environment described here is the same as shown in figure Figure 13-1 on page 310). This topology is as follows:

- VEWIN22 is the physical VMware VirtualCenter Server that gets all the management capabilities in its domain.

  Within VMware VirtualCenter you have two organizational objects: Farm Groups and Farms. The only purpose for Farm Groups objects is to group Farms to logical units. There is no equivalent mapping for Farm Groups through VMM into IBM Director.
- VE Farm is a logical organizational object that maps the VMware VirtualCenter object Farm onto IBM Director.

  The management capabilities for farm objects are in VMware VirtualCenter and all the modifications have to take place into the VMware VirtualCenter.
- MSNESX1 is the physical VMware ESX Server that is member of the farm VE Farm on the VMware VirtualCenter management domain VEWIN22. On this server two virtual machines are installed.
- The virtual machines (logical hardware platform) VM 1 and VM 2 on the VMware ESX Server MSNESX1 are, from an IBM Director standpoint ,a logical platform and have therefore the IBM Director label (in the Groups column of IBM Director is a group named Logical Platforms, open this group and you get in the Group Contents column all the virtual machines displayed that are known by the IBM Director (see Figure 13-2 on page 314).

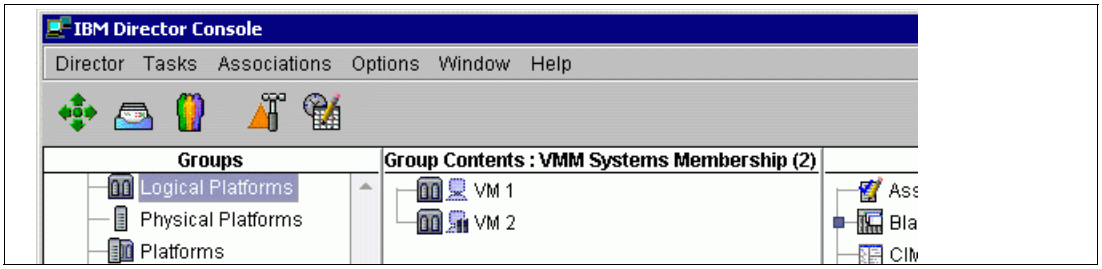

Figure 13-2 IBM Director Groups displays all Logical Platforms in the management domain

- On the virtual machines (logical hardware platform) VM 1 and VM 2 a Windows 2003 Server Standard Edition operating system image is installed. This operating system, named MSNVMG1, acts as a physical instance that is maintained through IBM Director. The IBM Director Agent is installed on this operating system and therefore the IBM Director Server has the full functional control over this machine from an IBM Director perspective.

Important: In the following, we name the guest operating system objects MSNVMG1 and MSNVMG2 running on their respective virtual machines VM1 and VM2. In VMware, Microsoft or others have the distinction between hardware and software in a management entity. IBM Director uses one object for both, the hardware and the software. It is one machine! Therefore, we use one common name for the hardware and software piece of two virtual objects: virtual machine. In cases where a distinction makes sense we notify it.

The following scenarios are a logical management flow through the IBM Director capabilities that makes usage of the VMM enhancements. We start at a manual level and then we raise to a more automatic (and so to speak autonomic) level. The main functions used are the power management facilities of virtual machines such as suspend and restart.

Note: Another powerful function that VMware provides through VMware VMotion is to move on the fly, while the users are working on a virtual machine, the image from one physical VMware ESX Server of another physical instance of VMware ESX server. In our environment we concentrate on the one VMware ESX Server scenario and remind the reader that it is an easy step to go from VMware power management functions to VMware VMotion functions. The only component that has to be changed is the action that is used in an event action plan.

13.3.2 VMM objects in IBM Director Console

To define VMM object in IBM Director Console, do the following:

1. Launch an action through the IBM Director Console. For example, initiate the power management resume task on a virtual machine. In our environment, right-click the **VM 2** object and choose **Power Management** → **Resume** to launch the resume task (see Figure 13-3).

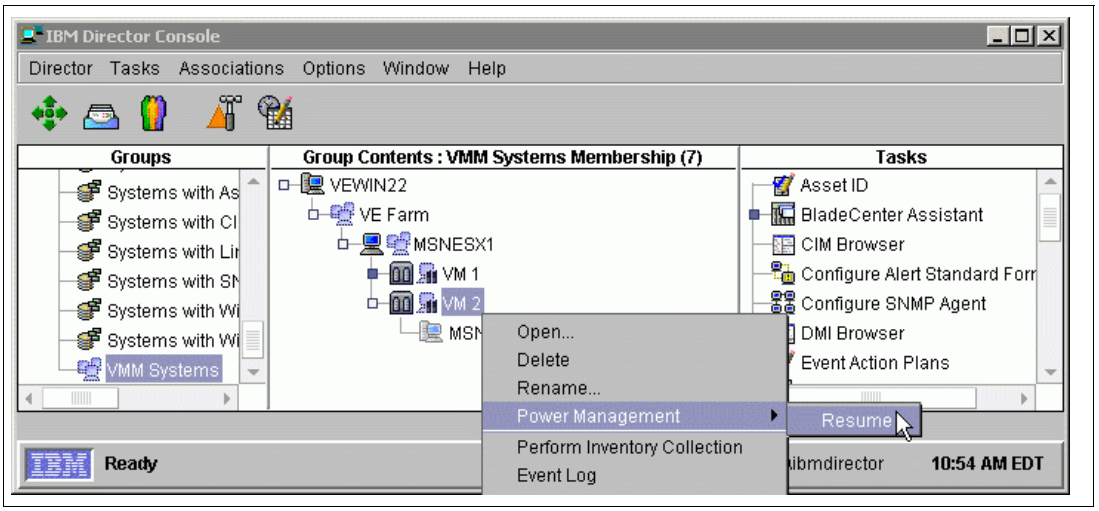

Figure 13-3 Launch Resume task through IBM Director Console

A window opens and asks if the execution should be scheduled or executed immediately. Choose execute immediately, and IBM Director displays the Execution History window. In this window, all status information such as process status, involved servers, and so forth for the launched command is displayed (as shown in Figure 13-4).

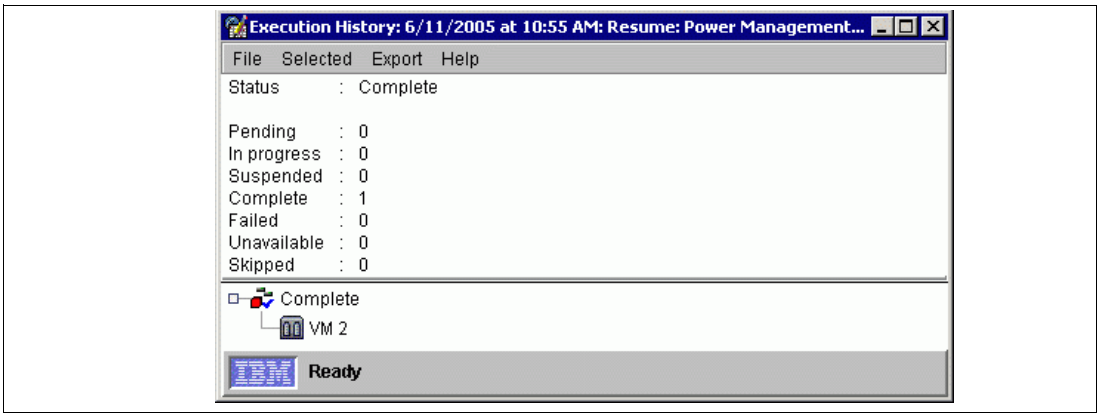

Figure 13-4 Execution history on VM 2 with Resume task

If the command is launched on a VMM object, the full execution log cannot be monitored through the IBM Director Execution window, because the execution is two-fold:

- The first part is the Web Service called from IBM Director to the coordinator object.
- The second part is the command execution in the coordinator object management domain.

Only part one, the status of the Web Service communication with the coordinator of the launched VMM object task, is displayed in the Execution History window. For the second

part the coordinator specific history and log mechanism must be used to get information about the second part execution.

2. To get more information about the different steps of the Web Service call, use the execution history log. In the Execution History window, select **File** → **View Log**. A new window lists all log entries. An example of a resume operation on VM 2 is shown in Example 13-1.

Example 13-1 Execution log on IBM Director Console for a VMM object **`Resume`** *command*

```
6/11/2005 10:55:01 AM    Job "Resume" activated.
6/11/2005 10:55:01 AM    Clients started for task "Resume"
6/11/2005 10:55:01 AM    Job activation status changed to "Active".
6/11/2005 10:55:02 AM    VM 2 client job status changed to "Active".
6/11/2005 10:55:02 AM    VM 2 - Power task successfull
6/11/2005 10:55:02 AM    VM 2 client job status changed to "Complete".
6/11/2005 10:55:02 AM    Job activation status changed to "Complete".
6/11/2005 10:55:02 AM    VM 2 - Power task successfull
6/11/2005 10:55:02 AM    VM 2 client job status changed to "Complete".
```

In the execution history, you can choose between three different log levels. Go to **View** → **Detail** in the Execution History window, and you have the choice between low, medium, and high log levels.

Consequently, the actual status representation of icons (such as the status icon on VM 2 in the transition from suspended to power-on) is not updated through IBM Director monitoring tasks. The update of the icons is pushed via Web Service calls from the coordinator to the IBM Director.

In our example, that means the icon representation of the VMM object virtual machine, VM 2, goes through Web Service updates from the coordinator as follows: suspended → transition → powered-on.

Figure 13-5 shows the VM 2 transition from powered-on to suspended from the top left to the bottom right. The reverse sequence shows the power management task suspend and its transitions on VM 2.

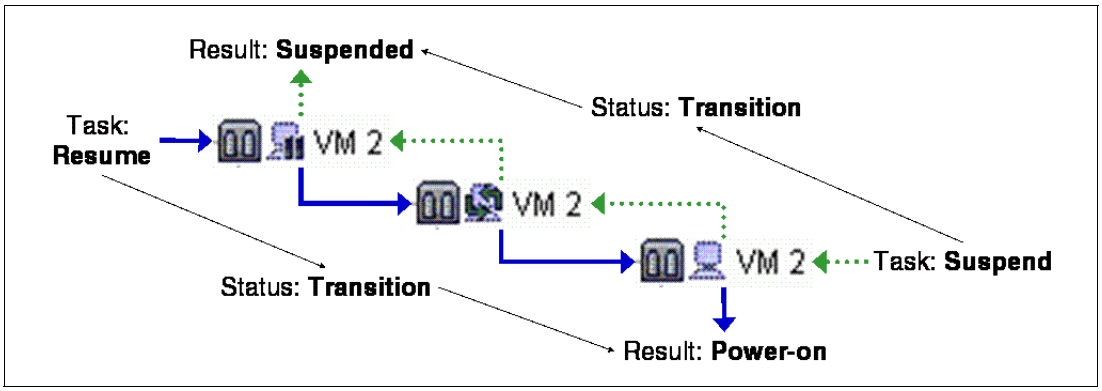

Figure 13-5 Task, transition, and result of IBM Director actions for virtual machine VM 2

Tip: The VMM objects in IBM Director console has a mouse-over function that displays the current state of that object. This is an easy way to get the most current status.

With this fundamental knowledge about the IBM Director console and the ability to start task and to monitor them in cooperation with the VMM objects, we can define the scenarios.

13.3.3 Running the scenarios

The following scenarios are build in a way to start from a more IBM Director Console interaction standpoint to a more automatic and self running way. Therefore we start simple in suspending a virtual machine through the Console and then introducing more the functions and features of the IBM Director server functions.

Note: In our scenarios the VMware function VMotion will not be considered. We concentrate in this redbook on one VMware ESX Server setup. VMware VMotion needs as minimum requirement two physical servers to move virtual machines from one server to another VMware ESX Server.

The VMware VMotion functionality is included through VMM in the IBM Director and the use of this functions are not different to the way we handle the power management functions in this redbook. Therefore it is very easy to enhance the scenarios here in only exchanging the action from power management to VMware VMotion actions.

For more information use the VMM and the corresponding VMware VMotion documentation available on the VMware Web site:

http://www.vmware.com/products/vmanage/vc_features.html

The following an overview of the scenario that we discuss in this chapter:

1. Suspend a virtual machine through IBM Director Console interface. This is a fast recapitulation of the sample used in the introduction of this chapter.
2. Define a scheduled job in IBM Director.
3. Scheduled tasks are an easy way to automate repeated functions.
4. Start a remote session on managed node.
5. With the IBM Director console you have a single point of access and management capabilities and with the remote session task we have a direct access to the operating system on a virtual machine through a command prompt.
6. Start a system account application to manage user accounts.
7. With the IBM Director Console you have a single point of access and management capabilities and with the system account task we have direct access to the account management on the operating system to work on users and group access management.
8. Start a resource monitor function on managed server.
9. We have the ability to monitor all resources IBM Director could query through VMM objects such as CPU utilization, disk usage, and many more.
10. Define a threshold through resource monitor.
11. Based on the resource monitor function, IBM Director defines the threshold that is monitored all the time. If a threshold limit is exceeded, events are launched that are used to launch actions.
12. Look into the event filter of IBM Director.
13. Understand the way IBM Director is using the event structure and how we could get the information and how to interpreted them.
14. Use Event Action Plan Builder to build launch automatically actions.

    This is the first piece that puts all the steps together to build one automated action. This action suspends VM 2 running an application App_HAL every time the application App_GNV has a CPU shortage on MSNVMG1 defined through the threshold setting.

The steps to implement this action are as follows:

a. Create action ACT Suspend VM 2.
b. Create event filter TEF MSNVMG1 CPU high.
c. Create event action plan EAP Suspend VM 2 on CPU shortage.
d. Connect event action plan to virtual machine MSNVMG1.
e. Test event action plan EAP Suspend VM 2 on CPU shortage.

15. Use command line tool genevent to launch IBM Director actions.

This is the second piece that introduces the IBM Director command **genevent** to send events to an event server. In our environment the IBM Director Server is named VEWIN21. The scenario that we use is that on MSNVMG1. The application App_GNV has finished its computation and sends to IBM Director an event to shut down and power off VM 1.

Suspending a virtual machine through IBM Director Console interface

To suspend a virtual machine through IBM Director Console interface, do the following:

1. Open the IBM Director Console.
2. Chose in the Groups column, **VMM Systems**.
3. Expand all objects in the Group Contents column (click the small quadrats in front of an object or use the arrow keys down and right to expand the tree).
4. Right-click the virtual machine that you want to suspend.

   Chose **Power Management** → **Suspend**. The sequence of the steps are displayed in Figure 13-6 for VM 1.

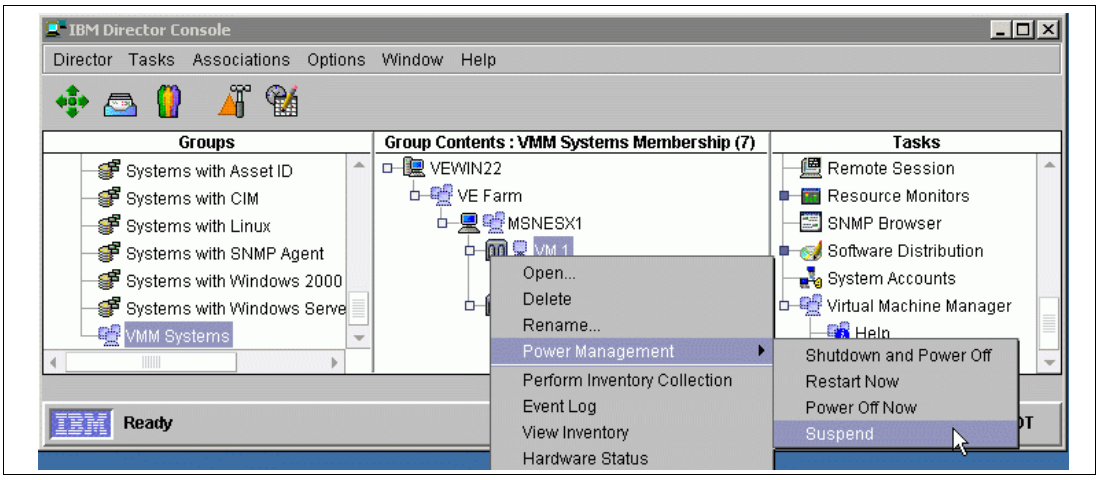

Figure 13-6 Suspend the virtual machine VM 1 through IBM Director console

A new window opens and asks for the user name and when the task should be executed. The options for executing the task are as follows:

- Schedule to start the task on a scheduled time. A new window pops up to define a new job name and all parameters needed to schedule a task.
- Execute Now to execute the task immediately after clicking that option
- Cancel to discard this task and do something else.

We choose Execute Now to run the task immediately. The Execution History window opens and the process status is in process, as shown in Figure 13-7 on page 319. If IBM Director can initiate the Web Service, the status changes immediately to complete. The coordinator

object (VMware VirtualCenter) forwards this request to the corresponding ESX Server. If the Web Service call fails, the Execution History window displays this status. You can look at the log for more information. For more information about the Execution History log file see at 13.3.2, “VMM objects in IBM Director Console” on page 315.

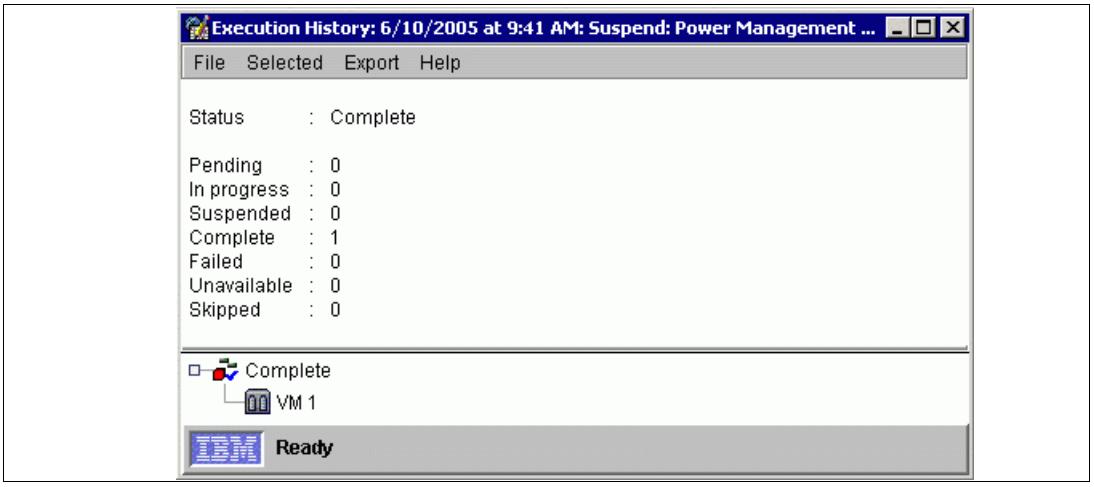

*Figure 13-7 IBM Director Console Execution History window for **VM 1***

The only way to watch the ongoing process is to look at the corresponding icon in the IBM Director Console. In our task, the icon representation switches from powered-on to transition to suspended. This is the lower-right to upper-left transition in Figure 13-5 on page 316 . You have to adopt VM 1 to match the figure.

Note: If a virtual machine is suspended (in our sample VM 1), the guest operating system that is running on that virtual machine (in our sample MSNVMG1) displays semi-transparently, and the IBM Director Agent functions are no longer available to this object until the virtual machine resumes.

After the virtual machine VM 1 is suspended you can right-click this object, chose **Power Management**, and the submenu offers the task Resume. If you want to resume that virtual machine, chose **Resume** with **Execute Now**, and the virtual machine resume.

If you need to execute a given sequence of tasks in a scheduled way (once or repeated), define the tasks with the IBM Director scheduled jobs. This is an automation to the manual way of executing single tasks on IBM Director objects through the IBM Director Console.

Defining a scheduled job in IBM Director

The reasons to create automated tasks are many fold. In the context of VMware with its virtual machines, we need to suspend low priority virtual machines in case that high priority virtual machines need more compute power. As an example, suspend an index service that mainly runs at night. Another scenario could be to copy an image at night time onto a backup location. With a scheduled task:

- We stop the image.
- We create the copy.
- We restart the image on a scheduled basis at night time.

In this section, we define a scheduled job for VM 2 to suspend the virtual machine at 10:05 a.m as follows:

1. On the IBM Director Console, launch the scheduler either by clicking the Scheduler icon on the IBM Director Console smart start bar or clicking **Tasks** → **Scheduler**. A new window Scheduler opens.
2. Choose the **Day Calendar** flag. In this view all scheduled jobs for today are displayed and sorted according to their schedules.
3. Create a new job by choosing **File** → **New Job**. A New Scheduled Job window opens, as shown in Figure 13-8.

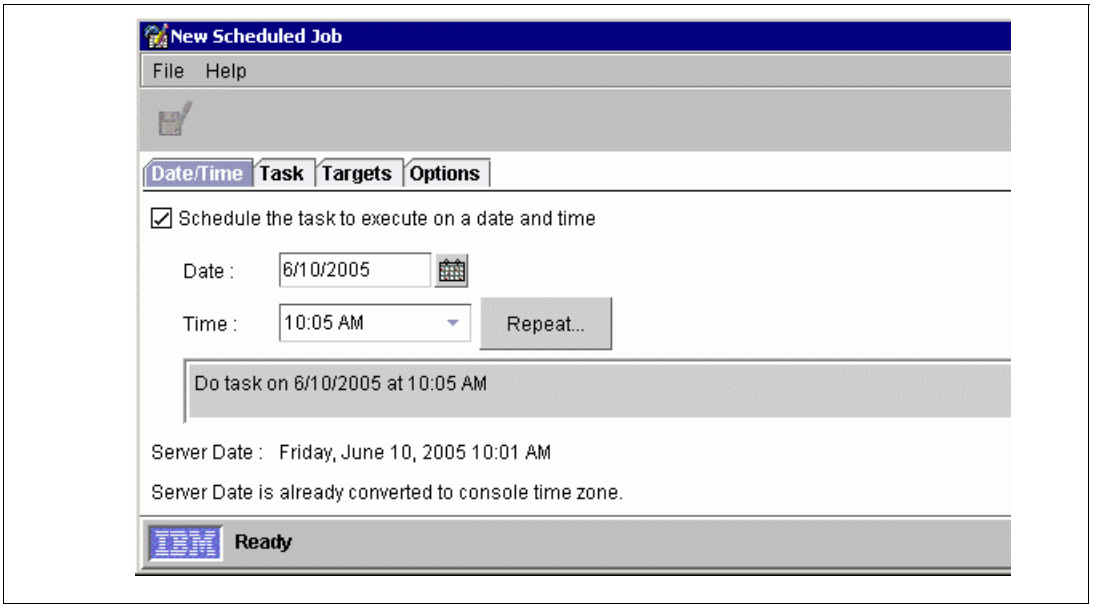
New Scheduled Job
File Help
Date/Time Task Targets Options
Schedule the task to execute on a date and time
Date : 6/10/2005
Time : 10:05 AM
Repeat...
Do task on 6/10/2005 at 10:05 AM
Server Date : Friday, June 10, 2005 10:01 AM
Server Date is already converted to console time zone.
Ready

Figure 13-8 Defining a new job schedule for VM 2 at 10:05 a.m.

4. Define the time that the job should run by entering `10:05 AM` in the Time field.
5. Click the **Task** tab to define the job that has to be executed at 10:05 a.m. through IBM Director.
6. Expand the Power Management task group to get all available power management tasks.
7. Click **Suspend** and either drag-and-drop the task into the Selected Task column on the right or click **Select**. The power management task suspend moves to the Selected Task column.
8. Click the **Targets** flag to define the server that has to execute the suspend task through IBM Director.
9. Look in the Available column for the target system, and click the object VM 2. Either drag-and-drop the object into the Selected column on the right or click **Add**. The target system **VM 2** moves to the Selected column.

Note: IBM Director supports a large variety of parameters and constraints on the job definition. We touch here only a basic scenario. For more information about how to create more sophisticated jobs, consult the IBM Director User's Guide at:

`http://ats.greenock.uk.ibm.com/Blade/Director/dir411_install.pdf`

10. Click **File** → **Save As** to save the new job description. Enter the name that defines the new job (such as `Suspend VM 2 @ 10:05 AM`) and then click **OK**.

The new job is created and moves to the Scheduler window. Until the job is waiting for execution, the icon for that job remains the same. After execution is completed, the icon changes to indicate the launch of the job.

11. To get information about the execution of the job, right-click the specific job in the Scheduler window and choose either **Open Execution History...** (the information is displayed as shown in Figure 13-7 on page 319) or choose **View Log...** (the information is displayed similar to that shown in Example 13-1 on page 316).

Starting a remote session on managed server

IBM Director does not differentiate between a physical server running IBM Director Agent or a virtual machine running a guest operating system on which the IBM Director Agent is installed.

With the remote session task, we start a remote session on a remote server by doing the following:

1. In our environment, we drag the object Remote Session onto the guest operating object MSNVMG1. In Figure 13-9, the changed cursor is shown for a drag-and-drop operation.

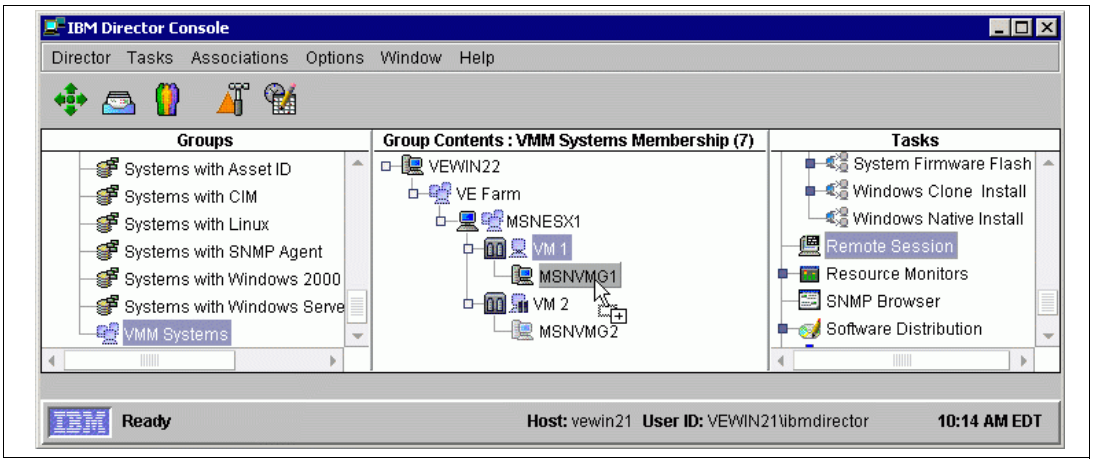

Figure 13-9 Drag the IBM Director object Remote Session and drop it on MSNVMW-1

2. A new Remote Session: MSNVMG1 window opens, and you have direct access on the operating system level. Thus, we have installed a Windows operating system on the virtual machine with a the remote session is a Windows command prompt starting on c:\. On a Linux operating system it could be for instance a Korn shell (depending on the user account default settings).

Now, you can start all command line commands on the remote session as though you are sitting in front of the server screen and keyboard (see Figure 13-10).

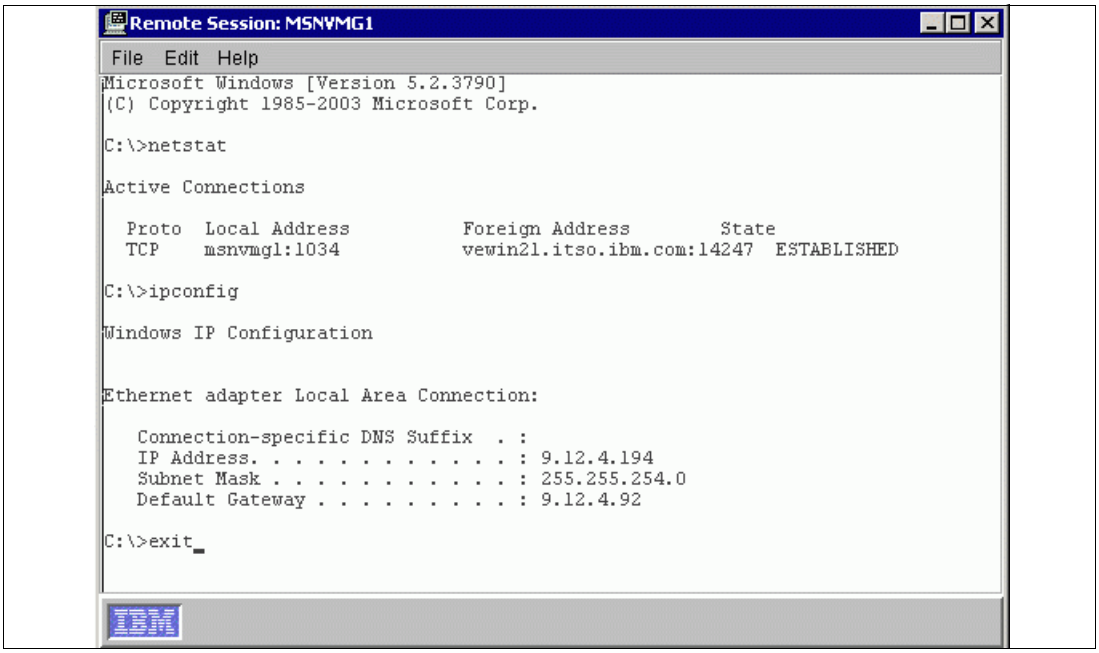
Remote Session: MSNVMG1
File Edit Help
Microsoft Windows [Version 5.2.3790]
(C) Copyright 1985-2003 Microsoft Corp.
C:\>netstat
Active Connections
Proto Local Address Foreign Address State
TCP msnvmg1:1034 vewin21.itso.ibm.com:14247 ESTABLISHED
C:\>ipconfig
Windows IP Configuration
Ethernet adapter Local Area Connection:
Connection-specific DNS Suffix . :
IP Address. . . . . : 9.12.4.194
Subnet Mask . . . . . : 255.255.254.0
Default Gateway . . . . . : 9.12.4.92
C:\>exit

Figure 13-10 Remote session Windows command prompt on MSNVMG1

3. To finish the work with the remote session, enter the **`exit`** command and the Remote Session window closes with the following message: `The remote session service has been closed`**.** Alternatively, you can use **File** → **Close** to close the remote session.

Note: On all servers in a IBM Director managed domain on which IBM Director Agents are installed, you could use the Remote Session task. Therefore, we can launch a Linux Bourne Again shell (**`bash`**) on MSNESX1 in our environment or a Windows command prompt on VEWIN22.

Starting a system account application to manage user accounts

With the IBM Director task **System Account** we manage the system account information about each server on which the IBM Director Agent is installed using IBM Director Console. An application is launched that allows the administrator to add, delete or update the account information about the managed server, on both object classes: users and groups.

- The way the system account task is started on a chosen server is the same way as described in “Starting a remote session on managed server” on page 321. You drag the task System Accounts and drop it on the corresponding object, such as MSNVMG1.
- A new window System Accounts: MSNVMG1 is displayed and we start administer the Users or Groups on that server. For more information go to the IBM Director Console User’s manual (see Figure 13-11 on page 323).

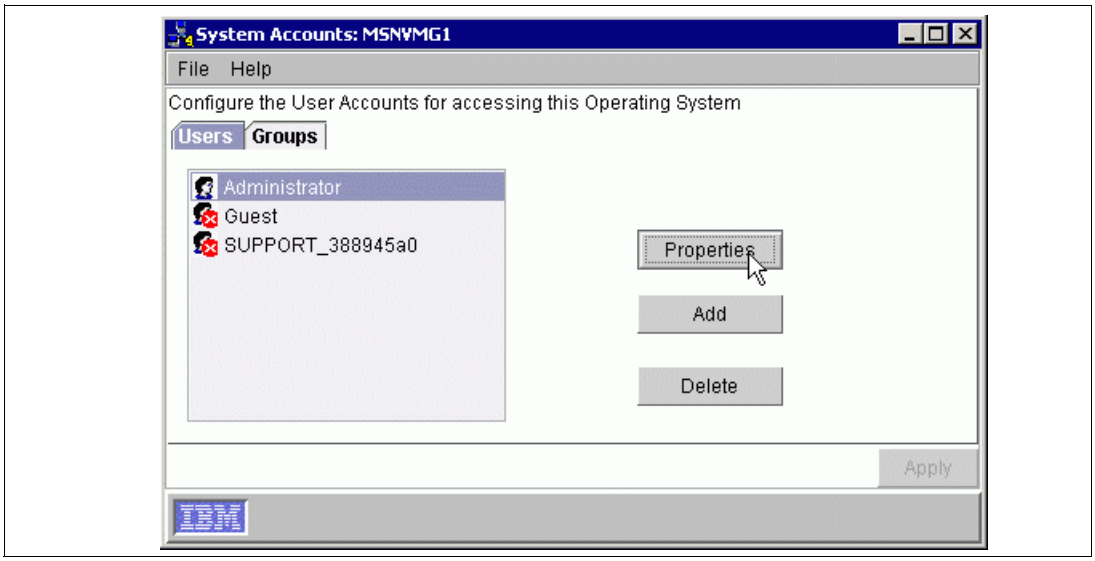

Figure 13-11 Manage system account information about MSNVMG1

Starting a resource monitor function on managed server

The IBM Director allows you to monitor all possible aspects of the hardware components and to send that information back to IBM Director Server in a vector of value pairs.

The way the resource monitor task is started on a chosen server is the same way as described in “Starting a remote session on managed server” on page 321. You drag the task Resource Monitors and drop it on the corresponding object, such as MSNVMG1. In this case a new Windows Resource Monitors (MSNVMG1) opens, and we select the monitors that we want to establish with the goal to monitor those resources. You start a resource monitor function by following these steps:

1. Click the small filled quadrant of the IBM Director Agent row in the Available Resources column.
2. Select **CPU Monitors**.
3. Drag the CPU Utilization monitor and drop it into a free space in the Selected Resources column on the right.

   IBM Director server starts immediately to collect the initial data from the resource and displays the message `Collecting Initial Data` in the newly created row in the Selected Resources column. After IBM Director has queried the information, the real value in a percentage is displayed in the [CPU Utilization] row.

Figure 13-12 shows the result of the description on how to define a monitor for a resource with the actual CPU utilization of 11%.

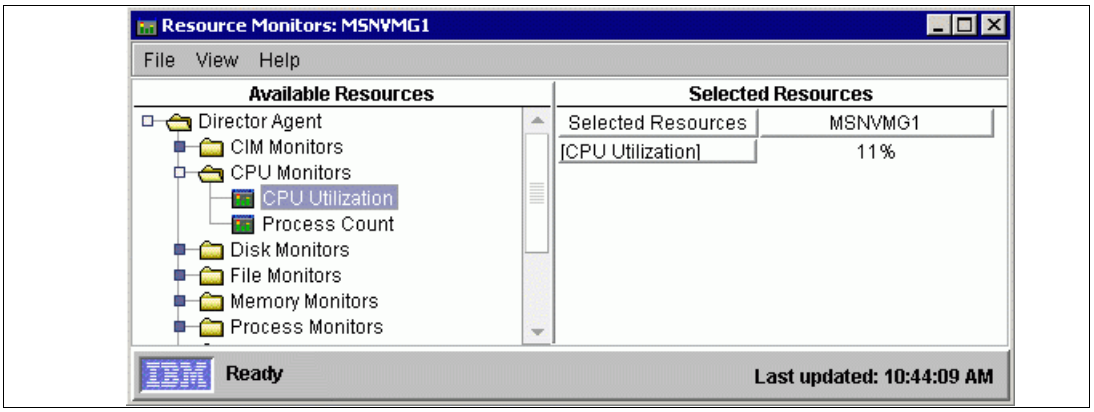

Figure 13-12 IBM Director Resource Monitor for the CPU utilization on MSNVMW-1

4. Click **File** → **Save As** To save the new created resource monitor.
5. Enter a meaningful monitor name. In our case, we use `Single CPU Utilization`. This new monitor is also added to the Tasks column beneath Resource Monitors.

Note: If we define a new monitor through the Resource Monitors main task, IBM Director opens a new window that contains no predefined resources. The Selected Resources column is empty. However, if we want to start with a predefined set of resources, we can expand Resource Monitors (click the small filled quadrat in front of Resource Monitors), and IBM Director lists the predefined monitors. If we choose, for example, the System Monitor monitor, the monitoring on that server starts with a set of monitors such as processor utilization, process count, drive space, and so forth.

Defining a threshold through a resource monitor

In this section, we define a threshold on a resource that then generates events in all cases where the defined level is exceeded or fallen below a denied limit. For a detailed description about the abilities and parameters of the usage of thresholds, we recommend that you consult the IBM Director User's Guide.

To define a threshold through a resource monitor, follow these steps:

1. In the IBM Director Console, drag the monitor from the Tasks column and drop it on the server MSNVMG1. A new Single CPU Utilization: MSNVMG1 window opens.

   IBM Director collects all values that are defined as resources for that monitor. It is recommended that you wait until the process is finished to be sure that all data needed on which to set a threshold is able to be monitored through IBM Director (see for reference Figure 13-12).
2. In the Selected Resources column, right-click the percentage value of the resource [CPU Utilization].
3. Choose **Individual Threshold...** to define CPU utilization thresholds for that resource on the server MSNVMG1. A new System Threshold: MSNVMG1 window opens.

   In the definition of a threshold, IBM Director offers a large variety of parameters. We have to define the labels, such as what event name is send in case a threshold is fired, and the values that defines that specific threshold, such as at which CPU utilization level is a High Error condition reached.

We use the following values to define our threshold on MSNVMG1:

- **Name** field: `ST MSNVMG1 CPU high`.
- **Description** field: `System Threshold on MSNVMG1 CPU utilization higher 95 percent for more than 15 seconds.`
- Select **Enabled to generate events**.
- Minimum Duration define 15 second(s).
- In the first Above Or Equal line, enter `95`. This is interpreted as a percentage value which defines the threshold High Error message.
- Click **OK** to save the new threshold.

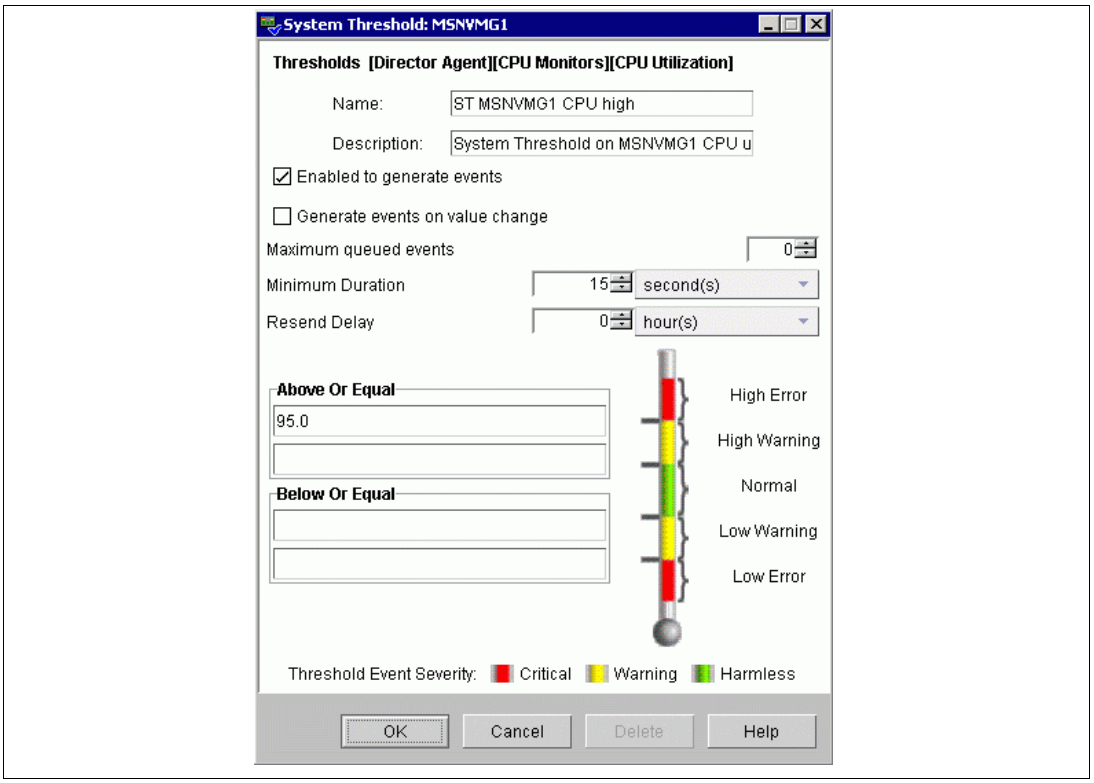

Figure 13-13 Define the CPU utilization threshold for MSNVMW-1

We have now created a new IBM Director monitor that keeps monitoring on MSNVMG1 according to the definition in Figure 13-13. This monitor now sends an event to IBM Director Server in every case that the CPU utilization is above or equal 95% for more than 15 seconds.

The 15 seconds duration is, for a real-life scenario, potentially too short. A realistic value might be a five-minute time frame. However, this is dependent on your purpose for monitoring and on the definitions of the event action plans.

The IBM Director Server tracks of all thresholds that are defined within an IBM Director domain. To look at all defined thresholds for a domain, go to the IBM Director Console Tasks column and open the hierarchy Resource Monitors. Right-click **All Available Thresholds** and select **Open**. A new All Available Thresholds window opens, and the foremost newly created threshold, `ST MSNVMG1 CPU high`, is in the list.

The second way to get a visual indicator of the existence of a threshold is to open a monitor. In our exercise “Starting a resource monitor function on managed server” on page 323, the [CPU Utilization] monitor has a small icon in front of the actual percentage value (a stylized

window). This is the indicator for a defined threshold on this resource. If the threshold is not reached, than the icon has a small blue check, if the threshold is fired, the blue check is replaced by a orange triangle. For more background information about the icons, go the IBM Director User's Guide.

Another way to track thresholds and the exceptions that are defined in a threshold is to monitor the IBM Director Console. However, this approach is only for a single minimal amount of thresholds that are feasible. A more convenient way is to look through the event filters of IBM Director into the system behaviors.

Looking into the event filter of IBM Director

All events that are generated throughout the system are stored and administered on the IBM Director Server. If you plan to make use of created thresholds, find the occurrence of such an event trough the Event Log in the Tasks column of the IBM Director Console. To look only in general at the event log, you could launch it by double-clicking the Event Log task. A new Event Log window opens, and all occurring events are listed there. (The default is the last 100 events.)

In our case, we are interested in the threshold violations on the CPU utilization of 95% over 15 seconds on MSNVMG1. The name of the threshold we defined is ST MSNVMG1 CPU high as described in "Defining a threshold through a resource monitor" on page 324.

To look into the event filter of IBM Director, we do the following:

- We open a specific filter in the Event Log task list that is named Critical Events. It is also possible to search within the All Events in Events Log but choosing the specialized filter critical events we have a smaller choice of events in the browser. Right-click that Critical Events log, and a window Event Log (Critical Events) opens. If there is no entry so far or no entry for the server MSNVMG1 than the CPU utilization on system MSNVMG1 went never higher than 95% for more than 15 seconds. Therefore, you have either to start CPU intensive jobs that fulfill the specific system boundaries for the threshold definition or lower the value in the threshold definition to get first messages into the log.
- You could use Refresh in the Event Log (Critical Events) smart menu to get the view updated. Click one event in the Events column with the System Name MSNVMG1 to get the details displayed on the right hand side column. There are normally also other events that regard to our IBM Director management domain. To be sure it is the last generated event, look at the Thresholds Name in the right column and it must be the name ST MSNVMG1 CPU high.
- Now we use the export function to a text file to save the information of an event log entry. There are three different file formats available:
  - Text File (.TXT)
  - XML Document (.XML)
  - HTML Document (.HTML)

  To save the event information right-click the designated event in the Events column and choose in the context menu Export → Text File (.TXT). In the pop-up dialog, specify the file name and the location on your file system (see Figure 13-14 on page 327).

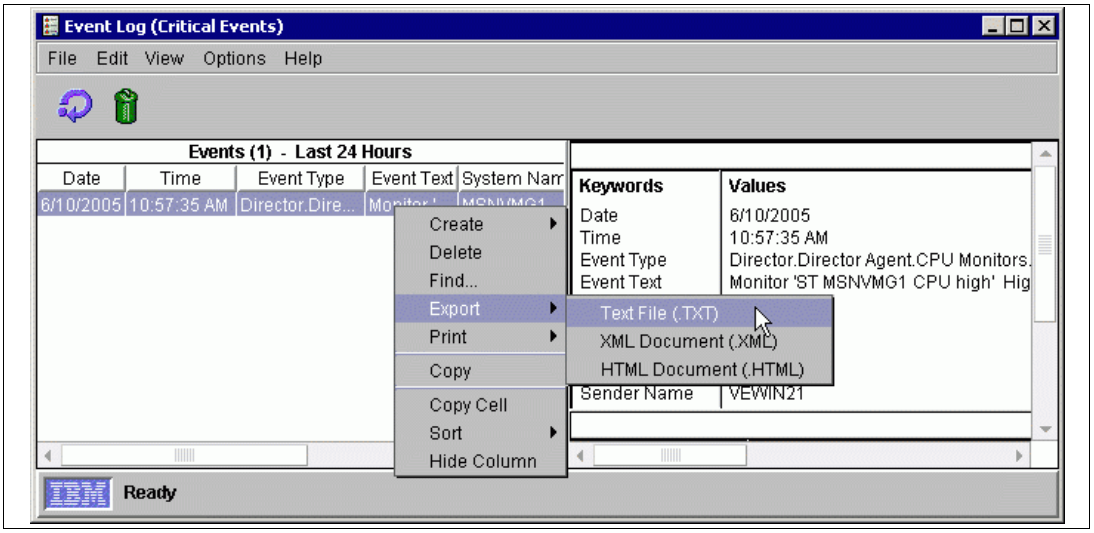

Figure 13-14 Save an event log entry to a text file

Example 13-2 shows the details of the system threshold that we defined on MSNVMG1.

Example 13-2 All data IBM Director has available on one event entry

```
TEXT="Monitor 'ST MSNVMG1 CPU high'  High Error: 'CPU Utilization' has been above or equal
to 95 for 0:00:18.  Value reported is 96.5.", TIMESTAMP=1118415455531, TIME="10:57 AM EDT",
DATE="6/10/2005", SEV=CRITICAL, CATEGORY=ALERT, SENDER="VEWIN21", GROUP="Not Applicable",
CORR="1:1118415455531", SNDUID="EADC7C2E4483A28F", SYSUID="37F6D31B064E9463",
TYPE="Director.Director Agent.MONCPU.2872344980.High Error", SYSTEM="MSNVMG1",
SLOT.ThresholdName="ST MSNVMG1 CPU high", SLOT.MonitorResource="CPU Utilization",
SLOT.ThresholdValue=95.0, SLOT.Duration=18, SLOT.ActualValueNumeric=96.49595642089844
```

In the next sections, we use thresholds and the mechanism of sending events to the IBM Director Server to do something valuable on the VMware environment. We first establish an action in the Event Action Plan Builder.

Using Event Action Plan Builder to build launch actions automatically

We have one VMware ESX Server with the following virtual machines running:

- One virtual machine is used to execute a production business application App_GLV on MSNVMG1
- One virtual machine is a long time running analyze tool App_HAL that posts data to a data mart on MSNVMG2.

The policy we have already defined is in case the CPU of MSNVMG1 is utilized above 95% for a period longer than 15 seconds an event ST MSNVMG1 CPU high is created and sent to the IBM Director (see also “Defining a threshold through a resource monitor” on page 324). We then create an event action plan based on that threshold that suspend the virtual machine VM 2 to release CPU cycles for MSNVMG1. To make the scenario more comprehensive, we create the following sequence of steps that could be reused at any time independently of each other:

1. Create action ACT Suspend VM 2.
2. Create event filter TEF MSNVMG1 CPU high.
3. Create event action plan EAP Suspend VM 2 on CPU shortage.
4. Connect event action plan to virtual machine MSNVMG1.
5. Test event action plan EAP Suspend VM 2 on CPU shortage.

Step 1: Creating action ACT Suspend VM 2

We first set up an automated action to define the action that has to take place. To define actions, start from the IBM Director Console and launch Event Action Plan Builder as follows:

1. Choose **Tasks** → **Event Action Plan Builder**. You can also start the builder through the smart icon task bar (the symbol is a hammer and an orange triangle). You can identify the corresponding Event Action Plan Builder icon by passing the cursor over the icons to get the mouse-over description. A new Event Action Plan Builder window displays as shown in Figure 13-15. The main columns of this window are:
   - A list of Event Actions Plans which are already implemented on this domain. There is a default implementation for All Events in this column.
   - A list of Event Filters based on the root from which source (duplication, exclusion, simple and threshold) they came.
   - A list of Actions that are available to be put into the event context.

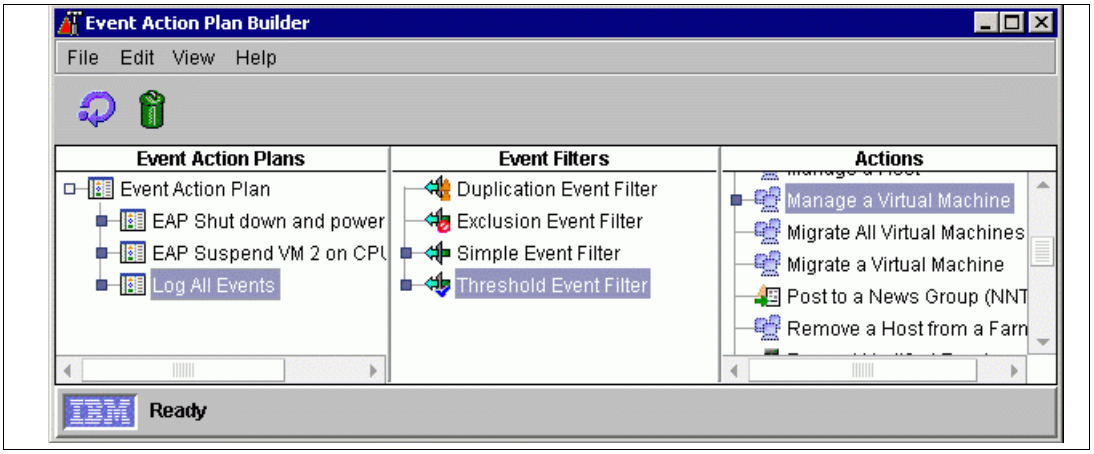

Figure 13-15 Event Action Plan Builder main window

2. Define in the Actions column a new action. This action is a Manage a Virtual Machine action that suspends VM 2.
3. Right-click **Manage a Virtual Machine** and in the Actions column.
4. Choose **Customize** and a new window Customize Action: Manage a Virtual Machine opens. In this window, you define a action that has to take place on a virtual machine. In our scenario, we chose **VM 2** and the action **Suspend** (see Figure 13-16 on page 329).
5. Choose **Save As** to save this action and enter `ACT Suspend VM 2` into the action name dialog.

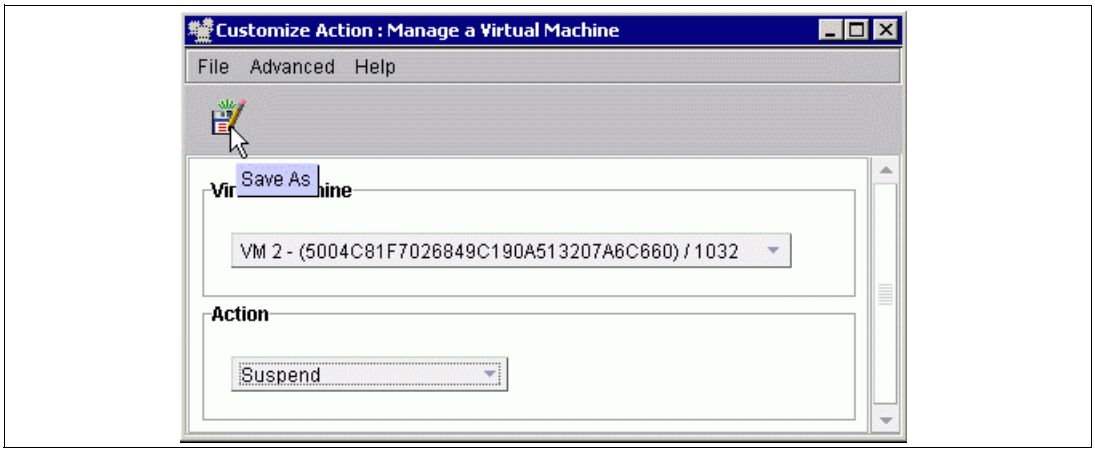

Figure 13-16 Definition of the action ACT Suspend VM 2 and start saving

Under Action, the newly created action displays as a sub-action beneath Manage a Virtual Machine.

Note: We plan to use the action history later. Therefore, we enable action history on the action ACT Suspend VM. Right-click the action and choose **Action History** → **Enable** to enable the history function.

6. Test the action Suspend VM 2. Right-click this action and choose **Test**. IBM Director launches the request instantaneously as defined. In our case the virtual machine VM 2 starts suspending. In the Event Action Plan Builder on the bottom of the window a message is displayed `Action has been launched, refer to action history for additional information`.
7. Right-click the action **ACT Suspend VM 2** and choose **Action History** → **Show**. A new window **Action History: Suspend Win 2003 Server SE 2** opens. In this window the details of the chosen action are displayed. If everything went well, the successful execution is reported. However, in case a specific action could not take place, the reason is listed. For example, if the action Suspend was chosen for a virtual machine and the virtual machine is already suspended, the error message Start Status will then be `Failed` and the Additional Information attribute will be `Target is in the wrong state for this operation`.

Now we have defined actions that could be launched in different ways such as schedule tasks and tasks launched through IBM Director Console. Another way is to define an Event Action Plan that includes the already threshold definition ST MSNVMG1 CPU high on MSNVMG1 and the action ACT Suspend VM 2. This enhancement is the focus of the next section.

Step 2: Creating event filter TEF MSNVMG1 CPU high

We create now a threshold event filter to filter only one specific event out of the masses of events posted to IBM Director by doing the following:

1. On the IBM Director Console in the Tasks column expand the category Event Log and double-click Critical Events. A new window Event Log (Critical Events) is displayed and in the Events column on the left all critical events are listed (for detailed information go to “Looking into the event filter of IBM Director” on page 326).

2. Choose one event that contains the text ST MSNVMG1 CPU high from MSNVMG1 and right-click the row. In the pop-up context menu, choose **Create** and click **Threshold Event Filter** (see Figure 13-17). A new window Threshold Event Filter Builder: New opens.

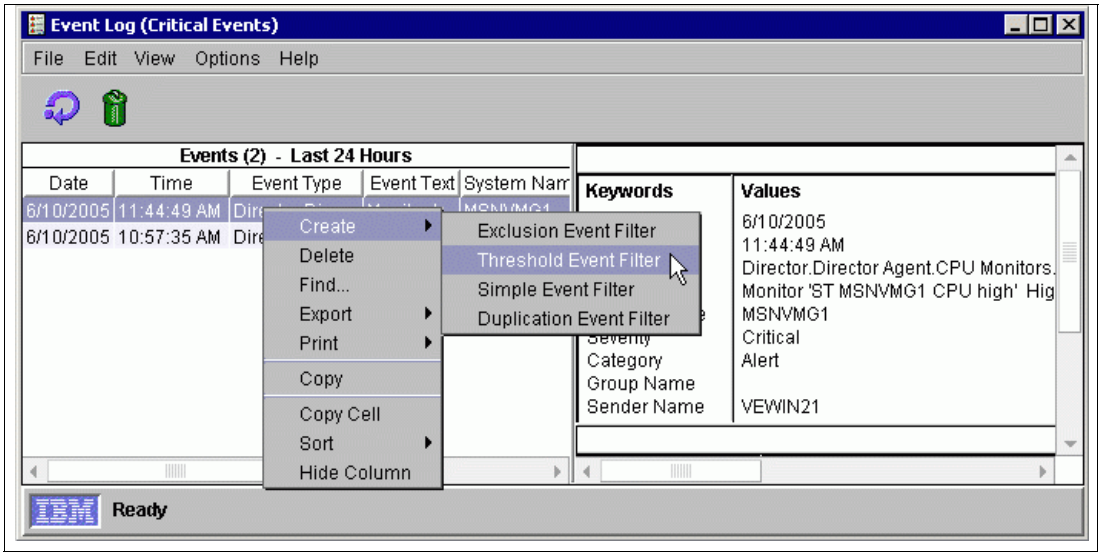

Figure 13-17 Create a new Threshold Event Filter through the Event Log

In the event filter builder all parameters are set to a value because we have launched the creation through an existing event therefore all of the parameters are initialized with the values of that chosen event. If you want to start from scratch you could do so in the Event Action Plan Builder window through the menu File → New → <...> Event Filter.

For verification purpose here is the list of the values

- Event Type check Director.Director Agent.CPU Utilization.High Error
- Severity check Critical
- Day/Time check Any
- Category check Alert
- Sender Name choose VEWIN21 (this is the server sending the message)
- Frequency set Interval to 0 second(s) and Count to 1

The only flag that need to be changed is Event Text because the transfer from a real event has put the whole event text into the text field and we define the threshold event filter more precise.

The original Event Text in the field is:

```
Monitor 'ST MSNVMG1 CPU high'  High Error: 'CPU Utilization' has been above or equal
to 95 for 0:00:19.  Value reported is 100.
```

3. Type in the text ST MSNVMG1 CPU high (see the definition we have used before in Figure 13-13 on page 325); now we have to use that exact definition to match the event with the filter) and choose the selection All words means all words must be included in the event text to get visible through that filter.
4. The last step is to save this new build threshold event filter. Click the icon to save or choose through the menu File → Save As and type a meaningful name in the pop-up Save Event Filter window, such as TEF MSNVMG1 CPU high.

Attention: We use a prefix to identify the object we are using. In this context we use **TEF** because it is the definition of a **T**hreshold **E**vent **F**ilter. Every time you see a given object you could easily identify from which object class the object is derived.

Now a new threshold event filter named TEF MSNVMG1 CPU high shows up in the Event Filters column of the Event Action Plan Builder window beneath the subgroup Threshold Event Filter.

We now have the threshold event filter created and could start putting all object together to build an event action plan.

Step 3: Creating event action plan EAP Suspend VM 2 on CPU shortage

All prerequisites are created for a new event action plan. We only have to create one new container that defines the new event action plan by doing the following:

1. Go to the Event Action Plan Builder window and right-click the row Event Action Plan in the Event Action Plans column. The context menu offers the only choice New and click it. A window Create Event Action Plan is displayed; type in the name for this new event action plan EAP Suspend VM 2 on CPU shortage. Click **OK** and the new event action plan show up in the Event Action Plan column.
2. Click the new created event action plan EAP Suspend VM 2 on CPU shortage and in the Event Filters column right-click TEF MSNVMG1 CPU high. A context menu displays and choose Add To Event Action Plan (see Figure 13-18).

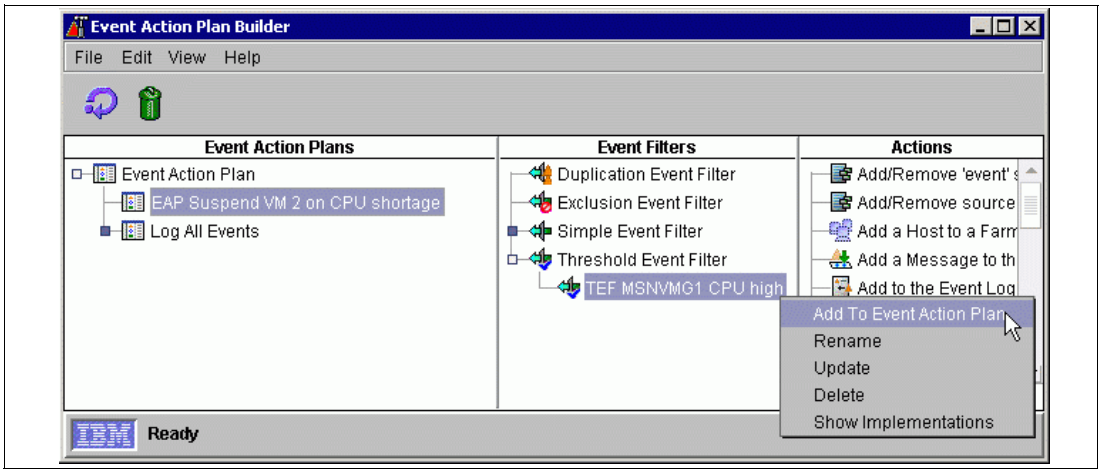

Figure 13-18 Add a Treshold Event Filter to an Event Action Plan

As a result of this action the threshold event filter shows up beneath our new event action plan EAP Suspend VM 2 on CPU shortage.

With this definition, the event action plan has a trigger that fires a defined event action plan that will be executed whenever a new event occurs that meets all criteria in that event filter. Differently worded: All events are filtered by the event filter TEF MSNVMG1 CPU high. Every time all criteria are met a new row is added to the filter with the event details. This adding a row than fire the event action plan EAP Suspend VM 2 on CPU shortage to start working. But what action needs to be executed?

We have to add an action to a event action plan.

3. Click the Event Action Plans column on the event filter TEF MSNVMG1 CPU high and right-click in the column Actions beneath Manage a Virtual Machine on the action ACT Suspend VM 2. A pop-up context menu is displayed. Choose **Add To Event Action Plan** and that action is put beneath the threshold event filter (see Figure 13-19).

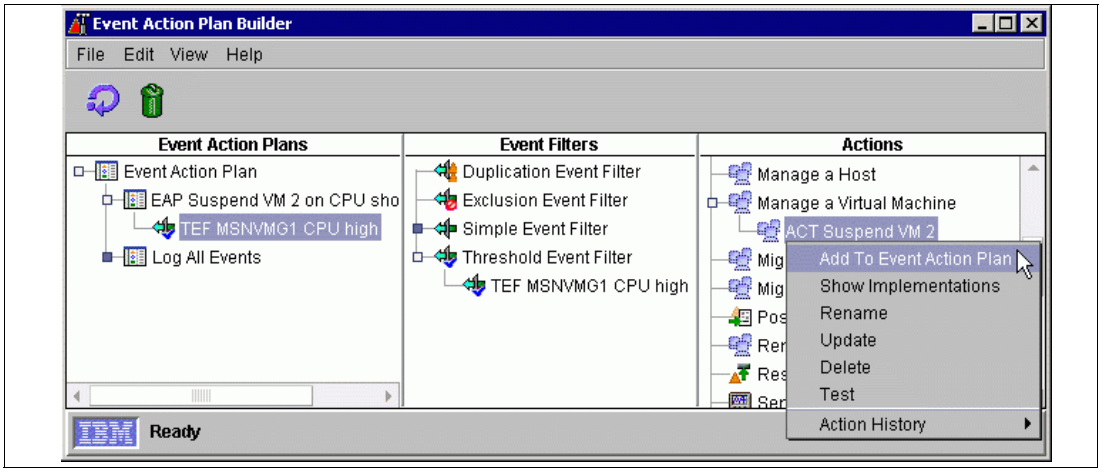

*Figure 13-19 Add an action **ACT Suspend VM 2** to the Event Action Plan Builder*

In the past steps, we defined an event action plan that is triggered through a threshold event filter. The filter launches an event every time the MSNVMG1 CPU is utilized over 95% for more than 15 seconds. The action that has to be taken is to suspend VM 2.

The last step is to connect the event action plan EAP Suspend VM 2 on CPU shortage to a IBM Director Server.

Step 4: Connecting the event action plan to virtual machine MSNVMG1

To connect the event action plan to IBM Director, we have to put the EAP Suspend VM 2 on CPU shortage event action plan onto the virtual machine MSNVMG1.

Remark: IBM Director generates only events if they are explicitly defined on the system (such as the event filter TEF MSNVMG1 CPU high). With event actions plans it is somehow the same: the definition itself affects nothing in the system. Only if we put an event action plan in context to groups or systems it could be executed.

To connect the event action plan, do the following:

4. On IBM Director Console expand in the Tasks column Event Action Plans and drag EAP Suspend VM 2 on CPU shortage and drop it on MSNVMG1 in the Croup Contents column. As a result of that action a pop-up window Information open and displays Event action plan has been added to selected group/system(s). Click **OK** to close the window.

   If the IBM Director Console has enabled the event action plans view (under Associations the property Event Action Plans is on) you see the added event action plan beneath MSNVMG1 in the IBM Director Console (see Figure 13-20 on page 333).

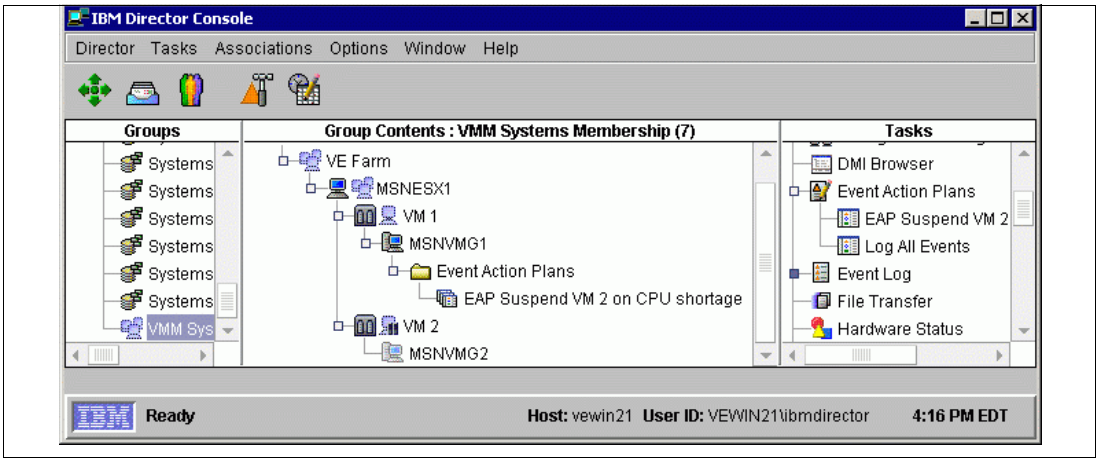

Figure 13-20 Event action plan EAP Suspend VM 2 on CPU shortage running on MSNVMG1

Everything is now set up andthe IBM Director started to look at the CPU utilization on MSNVMG1. In the next section, we test this EAP system.

Step 5: Testing the event action plan EAP Suspend VM 2 on CPU shortage

Here we look onto the same action as was described in “Using Event Action Plan Builder to build launch actions automatically” on page 327. Therefore, we open the action history through the Event Action Plan Builder window as follows:

1. In the Actions column, expand Manage a Virtual Machine and right-click **ACT Suspend VM 2**.
2. Choose in the context menu **Action History** → **Show**. A new window Action History: ACT Suspend VM 2 displays. The history might be not empty because the event of CPU shortage on MSNVMG1 has already taken place. For example, we used it before to test the action itself. Therefore, an entry log must be there. In that case, remove all events that start with an empty history.
3. As described in “Defining a threshold through a resource monitor” on page 324, we stress the virtual machine MSNVMG1 to get the threshold event fired. The threshold is exceeded only if the CPU is over 95% utilized for more than 15 seconds. The optical user interface representation is on the row of the MSNVMG1 icon at the time the threshold is reached. An icon is added to that row that has a small orange triangle, as shown in Figure 13-21 on page 334.
4. As a result of that threshold violation, the action ACT Suspend VM 2 is executed. Select Refresh the action history in the Action History: ACT Suspend VM 2 window if no action is listed.
5. Click in the Actions column on the most recent action and then, in the Action Details column, all details about the execution display. The End Status of this action is Successful. See also Figure 13-21 on page 334.

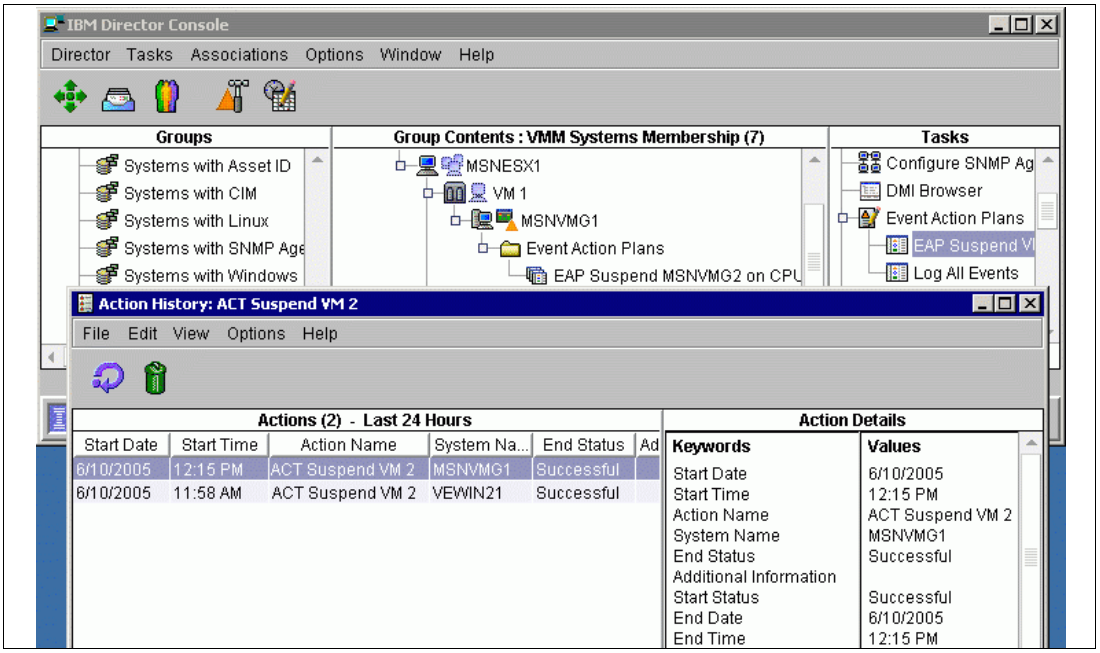

*Figure 13-21 Threshold reached on MSNVMG1 and action **ACT Suspend VM 2** is executed*

6. As final check, look at the VMware ESX Server and check if the virtual machine VM 2 is really suspending.

   Open on VEWIN22 the VMware VirtualCenter and click the Inventory flag. There, the object Discovered VMs in the left-hand side of the navigator. The state of VM 2 is `Suspending` and in the Active System Tasks on the Target VM 2 the action is `Suspend VM` (see Figure 13-22).

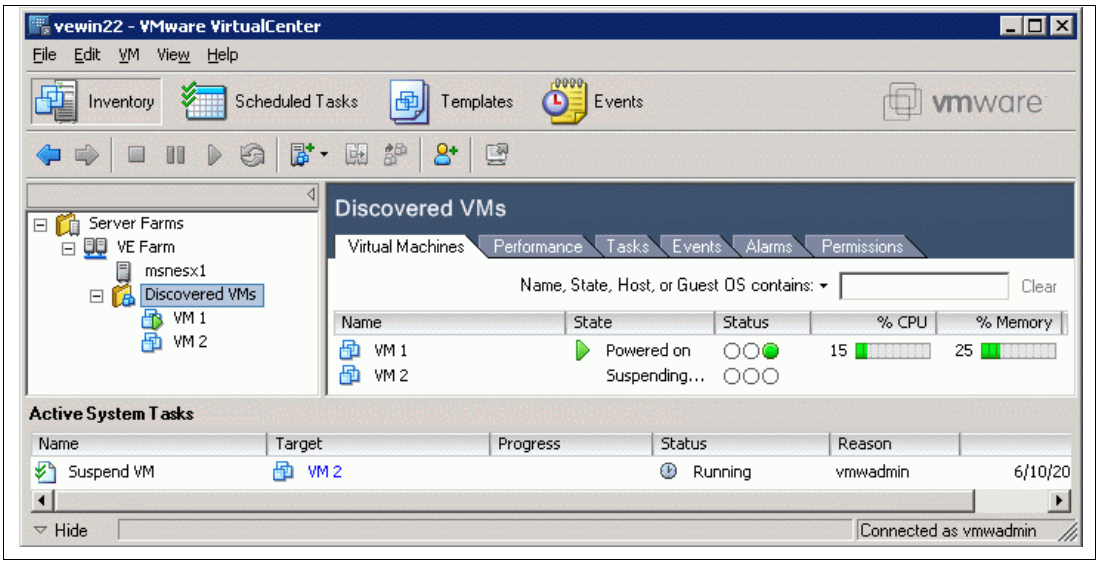

Figure 13-22 The event action plan suspending VM 2 is performed -- State: Suspending...

Another way to make use of the IBM Director actions is to use the **`genevent`** command to post events to an IBM Director Server. On this event, you could also define event filters for and build up event action plans to launch actions.

The next section introduces the **genevent** command. We reuse in most parts the things we have built up for the event action plan EAP Suspend VM 2 on CPU shortage.

Using genevent to launch IBM Director actions

In "Using Event Action Plan Builder to build launch actions automatically" on page 327, we defined an automatic action that is launched through a defined threshold monitor on the IBM Director. In this section, we use an event generating command line tool named **genevent** that is available on all servers on which the IBM Director Agent is installed.

We use this circumstance on the virtual machine MSNVMG1. The scenario is as follows:

- On the VMware ESX Server the virtual machine MSNVMG1 is used to execute a production business application App_GLV.
- This application runs only once a week for 10 contiguous hours.
- After the computation the virtual machine is no longer needed for 158 hours (7 days a week with 24 hours makes 168).

Because App_GLV knows when the computation for this week is done it could run a command after completion to send an event to the IBM Director to shut down and power off the virtual machine MSNVMG1.

Remark: There are many ways in the IBM Director tool set to reach the same goal to shut down a virtual machine. The intention we have here in mind is to show a new communication channel with **genevent** to the IBM Director and the manifold potentialities are only limited by our self imagination.

We reuse the sequence of the steps and definitions used in "Using Event Action Plan Builder to build launch actions automatically" on page 327 because this **genevent** scenario goes along the same line. Therefore, we point only to the corresponding section and formulate the deltas between there and here.

We follow the same sequence except we introduce the **genevent** generation:

1. Create action ACT Shut down and power off VM 1.
2. Use genevent the first time on MSNVMG1.
3. Create event filter SEF Shut down and power off VM 1.
4. Create event action plan EAP Shut down and power off VM 1.
5. Connect event action plan to virtual machine VM 1.
6. Test event action plan EAP Shut down and power off VM 1.

Step 1: Creating action ACT Shut down and power off VM 1

To start with this scenario, we create first the action that has to take place in case the event action plan is fired. Therefore we could use the information in "Step 1: Creating action ACT Suspend VM 2" on page 328 to create an action. The only changes we have are:

- The virtual machine we launch the action on is VM 1.
- The action we use is Shut down and power off.
- Save the new action with the descriptive event action name ACT Shut down and power off VM 1.

This action ACT `Shut down and power off VM 1 is` the action that the **genevent** launches through the event action plan. Therefore, the next step is to run the first time the **genevent** command on MSNVMG1.

Step 2: Using genevent the first time on MSNVMG1

We need to execute for the first time the **genevent** command on the command prompt on MSNVMG1. A convenient way is to use the IBM Director remote session function as described in “Starting a remote session on managed server” on page 321. Open a remote session on MSNVMG1 and enter the following command :

```
genevent /type:"Director.Director Agent" /text:"Shut down and power off VM 1"
```

The input and output on the remote session command prompt look similar to that described in Example 13-3.

Example 13-3 Execute **genevent** *command line tool on MSNVMG1*

```
C:\>genevent /type:"Director.Director Agent" /text:"Shut down and power off VM 1"
(C) Copyright IBM Corp. 1999, 2004 All Rights Reserved
The event was sent successfully to the given destination.
C:\>
```

Close the remote session on MSNVMG1 either by using the **exit** command or by using **File** → **Close**. A dialog window informs you that The remote session service has been closed. Click **OK** to close the dialog.

Note: The event type you are defining through the **genevent** is case-sensitive. If you write, for example, the **director** command with a lowercase letter, this is a different event type than if you use **Director** with an uppercase letter. You have to be aware of that when you search for event types.

Now a new event has been posted to IBM Director and we can look this event up in the event log to start creating a new event filter.

Step 3: Create event filter SEF Shut down and power off VM 1

The **genevent** command created a new event log entry in the event log we use to build a simple event filter from. Therefore, we could use the information in “Step 2: Creating event filter TEF MSNVMG1 CPU high” on page 329 to create an simple event filter. The only changes are:

- On the IBM Director Console in the Tasks column, expand the category Event Log and double-click All Events. A new window Event Log (All Events) is displayed.

  Note: The reason for not using the Critical Events is that the **genevent** command launches an event with the Severity `Unknown` by default. Under the Critical Events view, you only get events with Severity `Critical`, which means that the **genevent** generated event does not display in Critical Events.

- Search for one event in the Event Log (All Events) window that is from the system MSNVMG1. Select the event and look for the text that you used on the **genevent** command line. Right-click the row, choose **Create** → **Simple Event Filter**.
- In the Simple Event Filter Builder: New window verify the following values:
  - Event Type check Director.Director Agent
  - Severity check Unknown
  - Day/Time check Any
  - Category check Alert
  - Sender Name choose Any

Note: If we use the choice Any on Sender Name at this point we could also launch from any other system a genevent command to shut down this machine. If there is no need to have this flexibility then choose the machine name MSNVMG1

  - Event Text type `Shut down and power off VM 1` and choose the flag **All words**

- The simple event filters name is `SEF Shut down and power off VM 1`.

After we have saved the new event filter it shows up in the IBM Director column Tasks. Expand Event Log to see the filter SEF Shut down and power off VM 1.

Another place the new event filter is visible is in the Event Action Plan Builder in the column Event Filters, if the Event Action Plan Builder was open at the time generating the filter you first have to click the icon Update or go through the menu via View → Refresh. Expand Simple Event Filters to see the filter SEF Shut down and power off VM 1.

A new action is defined, the first event is generated on MSNVMG1 through the command genevent, a simple event filter is set. We now build up a new event action plan to combine all this parts.

Step 4: Creating event action plan EAP Shut down and power off VM 1

We reuse the information in "Step 3: Creating event action plan EAP Suspend VM 2 on CPU shortage" on page 331 to define the new event action plan. The only changes are:

- The name of the new event action plan is EAP `Shut down and power off VM 1`
- The event filter we use and add to the event action plan is SEF Shut down and power off VM 1
- The action we use and add to the event action plan is ACT Shut down and power off VM 1

The new event action plan EAP Shut down and power off VM 1 is defined and we have to connect this plan to the corresponding machine.

Step 5: Connecting event action plan to virtual machine MSNVMG1

The last step is to connect the event action plan EAP Shut down and power off VM 1 to a group or a system. See also "Step 4: Connecting the event action plan to virtual machine MSNVMG1" on page 332.

On IBM Director Console, expand in the Tasks column Event Action Plans and drag EAP Shut down and power off VM 1 and drop it on MSNVMG1 in the Croup Contents column. As a result of that action, a **Information** open and displays `Event action plan has been added to selected group/system(s).` Click **OK** to close this window.

Now we are ready to test the new established event action plan that has to be triggered through a genevent command.

Step 6: Testing event action plan EAP Shut down and power off VM 1

Some of the aspects we use in this section is already written down in "Step 5: Testing the event action plan EAP Suspend VM 2 on CPU shortage" on page 333. The steps there describe and how to verify the successfull execution of an event action plan if the trigger is a threshold.

In our scenario, the trigger is a command we have to execute on MSNVMG1. To do so, we use the method described in "Step 2: Using genevent the first time on MSNVMG1" on page 336 and launch the **`genevent`** command again.

As a direct reaction of the command a new event log entry is posted to the event log. To verify if this event has been posted go to the IBM Director Console, expand in the Tasks column Event Log and double-click All Events. In the list you find an entry in the Events column from the system MSNVMG1. Click that entry and the Event Text of the event is `Shut down and power off VM 1`.

The next verification step is to look at the action history. Open the Event Action Plan Builder go to the Actions column, expand Manage a Virtual Machine and right-click Shut down and power off VM 1. In the menu choose **Action History** → **Show**. A new window Action History: ACT Shut down and power off VM 1 open. Click the first entry in the Actions column and the Action Details on the right tells you on the keyword `End Status` that the action has been executed successfully.

As final check we have to look onto the VMware ESX Server and check if the virtual machine VM 1 is really suspending. Open on VEWIN22 the VMware VirtualCenter and click in the Inventory flag. There the object Discovered VMs in the left hand side navigator. The state of VM 2 is Stopping... and in the Active System Tasks on the Target VM 2 the action Power off VM is taken place (see Figure 13-23).

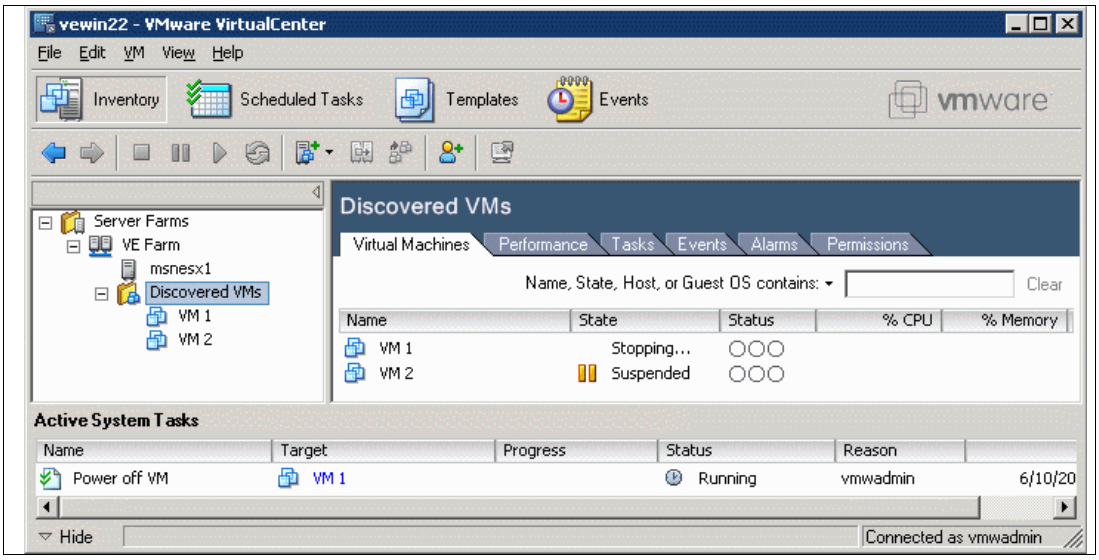

Figure 13-23 The event action plan shut down

13.4 Miscellaneous

In this section, we give some hints and tips for some situations that might occur in an environment with IBM Director and VMM.

VMware VirtualCenter credentials are broken

The communication between the IBM Director and the VMware VirtualCenter takes place through a Web Service. Therefore, there are two possible ways that might have problems in a IBM Director with VMM and VMware VirtualCenter communication:

- The Web Service as a transport mechanism for service calls. The problem could be solved either in restarting the VMware VirtualCenter Web Service through the Windows Service (see Figure 13-24 on page 339) or in restarting the whole VMware VirtualCenter Server service.

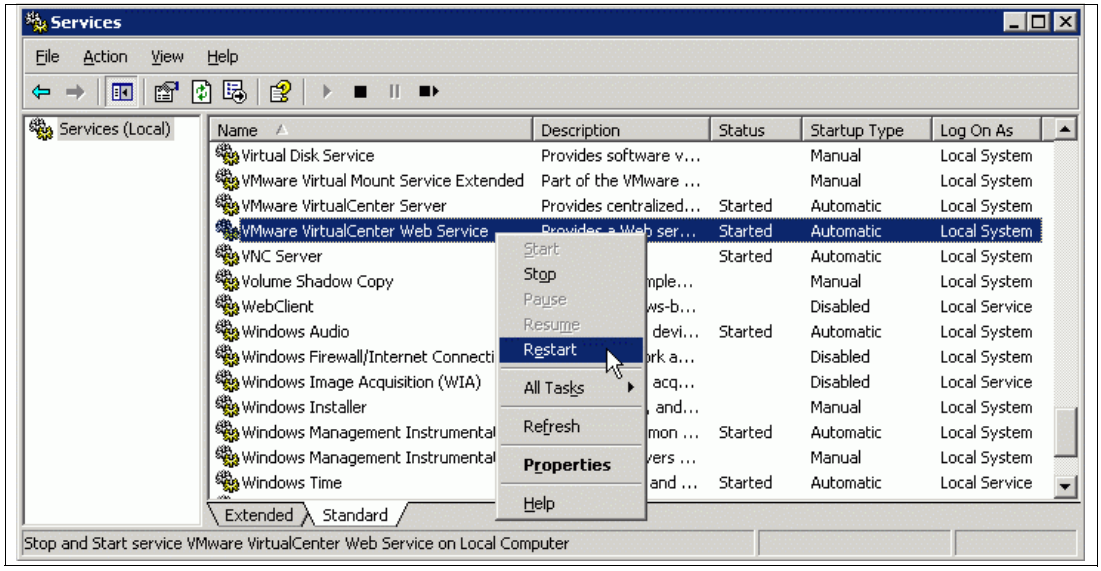
Services
File Action View Help
Services (Local)
Name Description Status Startup Type Log On As
Virtual Disk Service Provides software v... Manual Local System
VMware Virtual Mount Service Extended Part of the VMware ... Manual Local System
VMware VirtualCenter Server Provides centralized... Started Automatic Local System
VMware VirtualCenter Web Service Started Automatic Local System
VNC Server Started Automatic Local System
Volume Shadow Copy Manual Local System
WebClient Disabled Local Service
Windows Audio Started Automatic Local System
Windows Firewall/Internet Connecti Disabled Local System
Windows Image Acquisition (WIA) Disabled Local Service
Windows Installer Manual Local System
Windows Management Instrumental Started Automatic Local System
Windows Management Instrumental Manual Local System
Windows Time Started Automatic Local Service
Start
Stop
Pause
Resume
Restart
All Tasks
Refresh
Properties
Help
Extended Standard
Stop and Start service VMware VirtualCenter Web Service on Local Computer

Figure 13-24 Restarting the Web service on the VMware Virtual Center server

- The execution on the VMware VirtualCenter server as a component. For issues with the communication based on up and running Web Services, we revoke the credentials we have defined on the VMM coordinator object VEWIN22.

  Therefore, in our environment go to IBM Director Console and right-click VEWIN22 in the Group Contents column. In the pop-up context menu choose **Coordinator Management** → **Revoke Credentials** and the access to all VMM objects are revoked. You could only continue work on that machines that are natively integrated into IBM Director. In our environment all machines that has the IBM Director agents installed such as MSNVMW1, MSNVMG1 and MSNVMG2. On the VMM object for farms and hosts it is no longer supported.

  Then enter again the credentials for the coordinator object VEWIN22. Right-click VEWIN22 in the Group Contents column. In the pop-up context menu choose **Coordinator Management** → **Enter Credentials** and a window will open to enter the credentials. For details how to enter credentials look at “Verifying the VMM Agent installation” on page 234.

  If the answer of the window Enter Credentials is `Not Accepted!` the communication through Web services is established but the VMware VirtualCenter server does not reply to that request.

  Therefore, we eliminate the Web Service communication and start working on VEWIN22 itself. In an browser, enter the following:

  ```
  https://localhost:8443/login.html
  ```

  If you get the following in the browser then the account information about that VMware VirtualCenter is broken (see Figure 13-25 on page 340):

  ```
  <env:faultstring>Could not retrieve account information
  ```

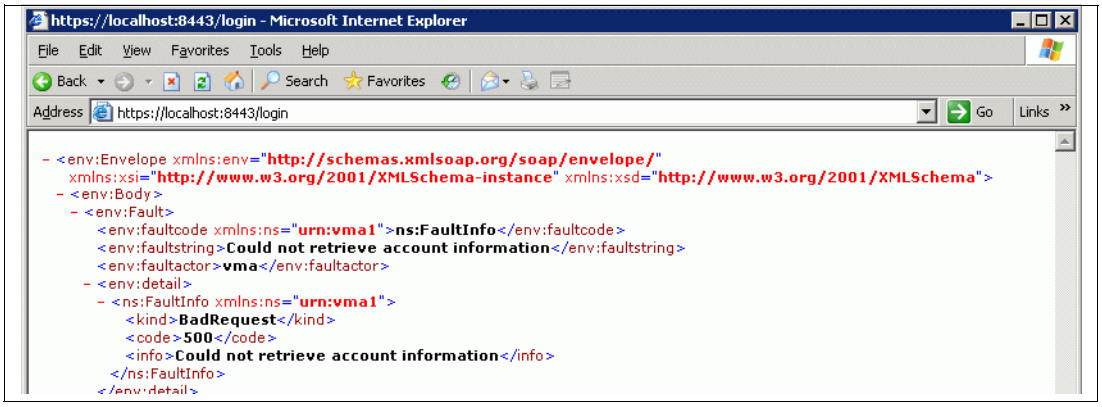

Figure 13-25 The VMware VirtualCenter account information are broken

To reset the account information to the value you have defined at the installation time you could use the **vma** command from VMware to do so:

- Search on the Windows system for the program vmw.exe. In our environment it resides on C:\Program Files\VMware\VMware VirtualCenter and execute the command that we have used it in Example 13-4:

```
vma -update -username vmwadmin -password password
```

Example 13-4 Command line call of vma.exe to set vmwadmin account information

```
C:\Program Files\VMware\VMware VirtualCenter>vma -update -username vmwadmin -password
password
[DATE 2005-06-02 18:49:20 verbose] Updating user information
[00:00:00.000 verbose] subject /vcenter username: vmwadmin
[00:00:00.000 verbose] subject /vcenter updated password.

C:\Program Files\VMware\VMware VirtualCenter>
```

If you now request the URL `https://localhost:8443/login.html` through a browser on the VWEWIN22 it allows you to get the definitions on that server in a XML notation VMware VirtualCenter is position the service definition in XML format.

For a complete test, see B.2.3, “Verifying VMware VirtualCenter server installation” on page 372.

Cloning a VMware Image through the VMware VirtualCenter functions.

If you clone a VMware image the cloned system will look similar to the original hence to the nature of cloning. Because this behavior brings some positive effects, such as faster availability of new like-systems, and some bad effects, such as system unique parameters are also cloned that has to be at any time unique.

In our context the uniqueness of a system is represented through attributes such as **System Name**, **Unique System ID**, **TCP/IP address**, **TCP/IP Host name**, **System UUID**, **MAC Address** to name only some of them. For example in a Windows environment most of the attributes could be changed either through updating the **Network Connection** of the network card (on Windows 2003 Server goto: **Start** → **Control Panel** → **Network Connections** and chose the corresponding network card) or through the System Properties of the Windows system (on Windows 2003 Server goto: **Start** → **Control Panel** → **System** and chose the flag **Computer Name**).

But before you could reboot the system, you have to remove the **Unique System ID** information. The **Unique System ID** will be created on system start if the information is missing. Start the Microsoft Registry Editor (for instance execute the command **regedit** on a command prompt) and search for the following location:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\ComputerName\ComputerName
```

Delete the attribute TWGMachineID from this location (see Figure 13-26).

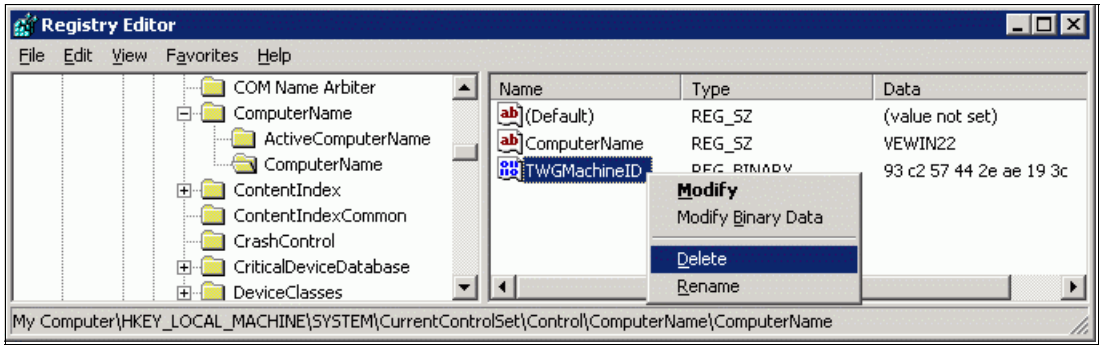

Figure 13-26 Delete TWGMachineID attribute in Windows system registry

Also go with the Windows Explorer to the directory:

C:\Program Files\IBM\Director\data\

Delete the file twgmach.id (see Figure 13-27).

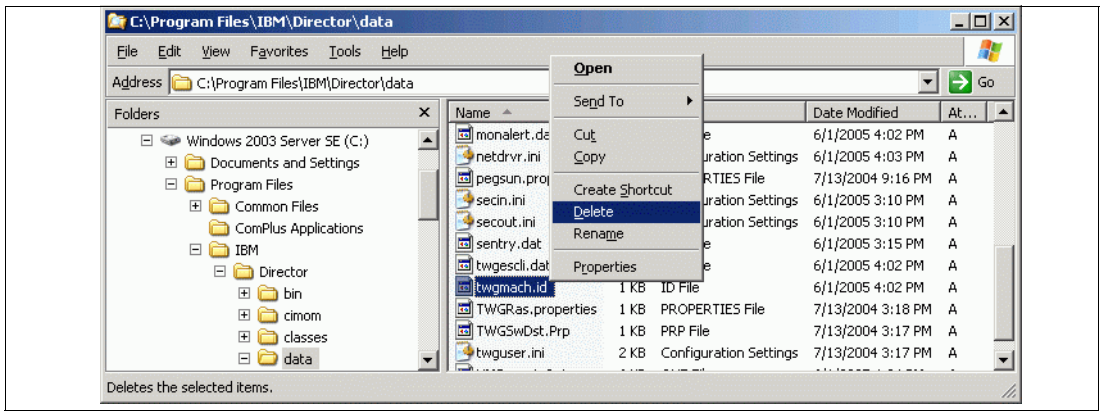

Figure 13-27 Delete the twgmachin.id file from IBM Director data directory

These two are responsible for storing the Unique System ID of a IBM Director Agent.

After rebooting the Windows system the representation of this computer will have different System Unique ID in the Display System Attributes window in the IBM Director Console. And will therefore be treated as a unique instance.

IBM Director action is not executed

In case IBM Director actions are not executed even thought the triggering event has taken place (look up in the event log that the event is there) one thing you have to look at is that the parameters for that action are anymore valid.

We assume that in our case the action ACT Suspend VM 2 has not been executed but the event log has the corresponding event that triggers the event action plan. If so, then do the following:

1. Go first to the Event Action Plan Builder and select in the Actions column the action **ACT Suspend VM 2**. Right-click that action and choose **Action History** → **Show**. In the new window **Action History: ACT Suspend VM 2** click the most actual event. In the **Action Details** column look at the **End Status** keyword. If it is `Not performed` than look further on the **Additional Information** keyword. In our case, the additional information that is given is `Failed to initiate operation -- invalid parameter`. That means something is wrong with the parameters of the action we are trying to launch (see Figure 13-28).

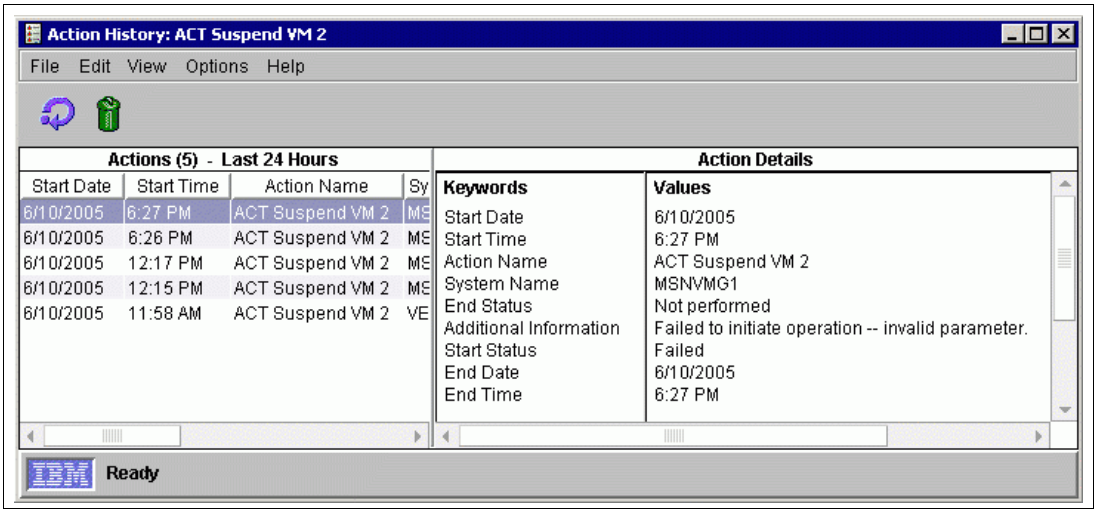
Action History: ACT Suspend VM 2
File Edit View Options Help
Actions (5) - Last 24 Hours
Start Date Start Time Action Name Sy
6/10/2005 6:27 PM ACT Suspend VM 2 MS
6/10/2005 6:26 PM ACT Suspend VM 2 MS
6/10/2005 12:17 PM ACT Suspend VM 2 MS
6/10/2005 12:15 PM ACT Suspend VM 2 MS
6/10/2005 11:58 AM ACT Suspend VM 2 VE
Action Details
Keywords Values
Start Date 6/10/2005
Start Time 6:27 PM
Action Name ACT Suspend VM 2
System Name MSNVMG1
End Status Not performed
Additional Information Failed to initiate operation -- invalid parameter.
Start Status Failed
End Date 6/10/2005
End Time 6:27 PM
IBM Ready

Figure 13-28 Status is Not performed because of invalid parameters

2. Go back to the Event Action Plan Builder and right-click the action ACT Suspend VM 2 in the Actions column. Click **Update**. The Customize Action: ACT Suspend VM 2 opens and the virtual machine is Previous Selection Invalid. Select **New Object** (see Figure 13-29).

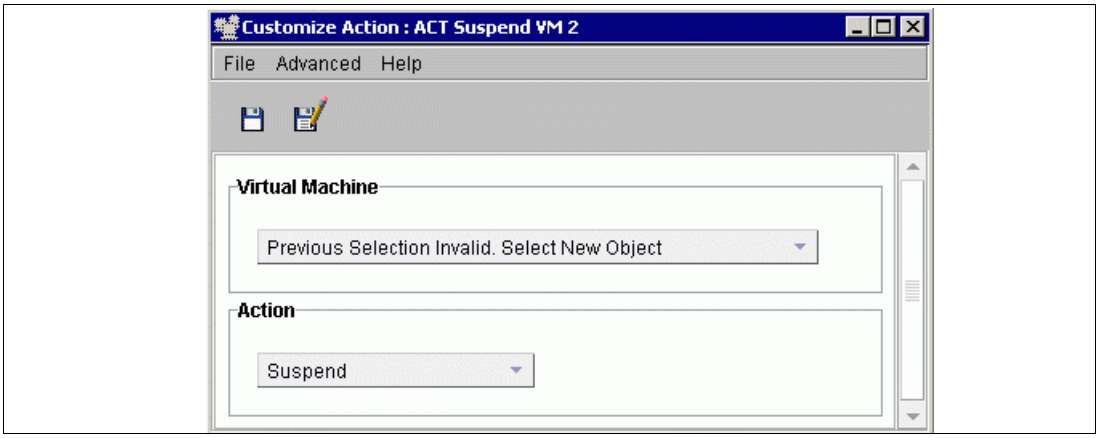

Figure 13-29 Virtual machine setting is invalid: Previous Selection Invalid. Select New Object

3. Go to the Virtual Machine section and choose on the pull-down menu the virtual machine VM 2 for that action. Save the action and try to run the scenario again. It should now work.

Note: One reason that IBM Director disconnects from a specific virtual machine is that the VMware ESX Server is not running all the time. IBM Director stops querying that resource and sets the actions based on that virtual machine, respectively, VMware ESX server into an invalid state. It is the administrators responsibility to relink the corresponding link after re-establishing the connection.

Part 4

Summary

In this short part, we summarize the success key points for a quick and successful Virtualization Engine installation and configuration.

The first thing to understand is the problem you want to solve. If the problem is not well understood or not well expressed, your customers might be disappointed by the solution.

The Virtualization Engine installation requires many components that need to be at the right level to be sure they communicate altogether. Understand that the Virtualization Engine is not complex by itself: what is complex is the problem it wants to fix. A heterogeneous environment, the environment where the Virtualization Engine adds most of its value, is already difficult to manage. The Virtualization Engine does exist to simplify this type of environment.

In each of the preceding chapters, we used the notes and tips shaded boxes that provide key points to consider when installing or using the Virtualization Engine components. They can help when installing each component.

For a successful Virtualization Engine implementation, the following sections are key.

Read Information Center News

The Virtualization Engine is made of multiple components, running on multiple operating systems. Today, only some of the components run on specific operating systems: on time more and more components will be made available. The Information Center News is the place to find the latest about operating systems support, fixes, and system requirements acquired from the latest experience. The Information Center News is available at:

http://publib.boulder.ibm.com/infocenter/eserver/v1r1/en_US/index.htm?info/veicinfo/eicarplangeneral.htm

Figure 1 on page 346 shows the What's New information and what you can gather from there.

© Copyright IBM Corp. 2005. All rights reserved.

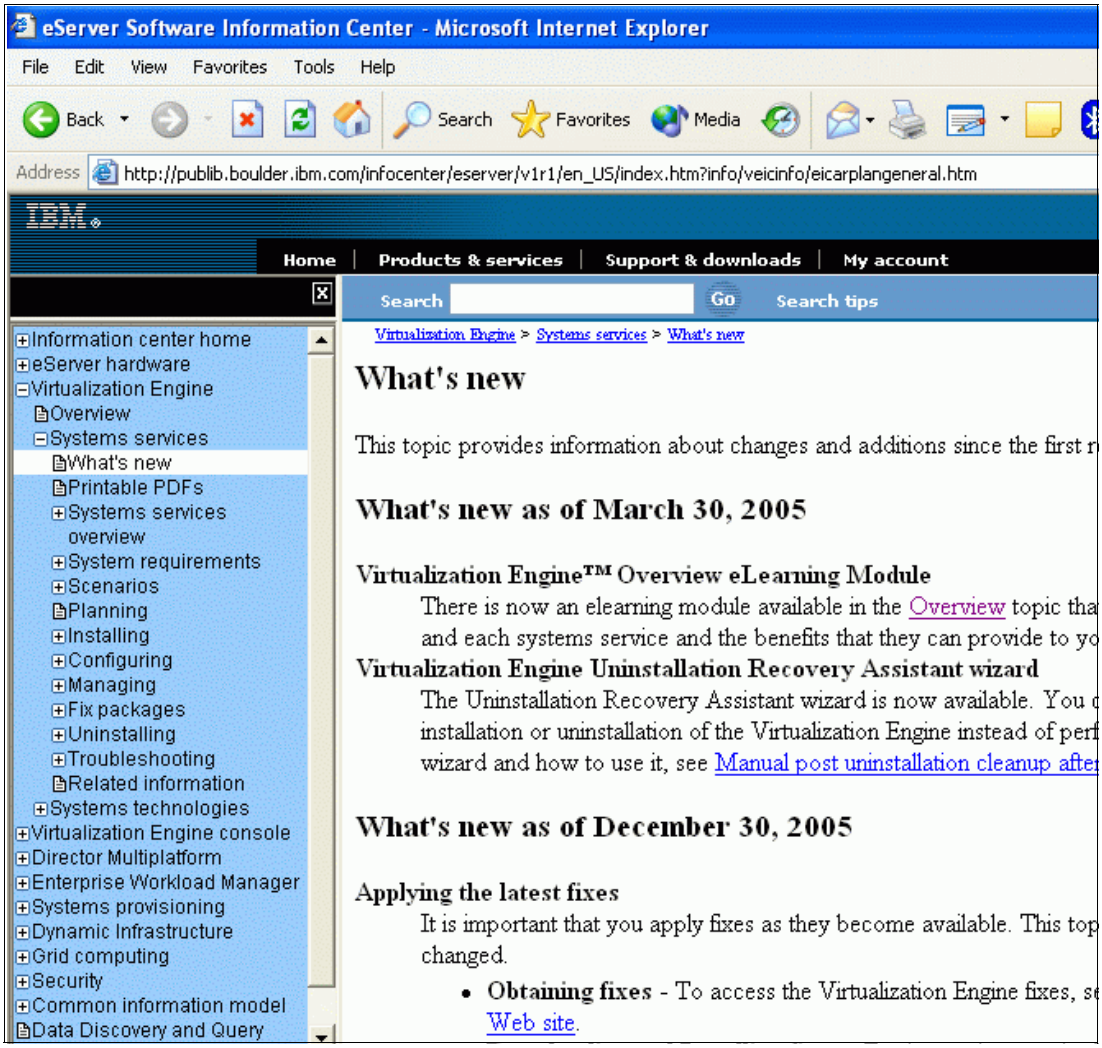
eServer Software Information Center - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address http://publib.boulder.ibm.com/infocenter/eserver/v1r1/en_US/index.htm?info/veicinfo/eicarplangeneral.htm
IBM
Home | Products & services | Support & downloads | My account
Search Go Search tips
Information center home
eServer hardware
Virtualization Engine
Overview
Systems services
What's new
Printable PDFs
Systems services overview
System requirements
Scenarios
Planning
Installing
Configuring
Managing
Fix packages
Uninstalling
Troubleshooting
Related information
Systems technologies
Virtualization Engine console
Director Multiplatform
Enterprise Workload Manager
Systems provisioning
Dynamic Infrastructure
Grid computing
Security
Common information model
Data Discovery and Query
Virtualization Engine > Systems services > What's new
What's new
This topic provides information about changes and additions since the first r
What's new as of March 30, 2005
Virtualization Engine™ Overview eLearning Module
There is now an elearning module available in the Overview topic tha
and each systems service and the benefits that they can provide to yo
Virtualization Engine Uninstallation Recovery Assistant wizard
The Uninstallation Recovery Assistant wizard is now available. You
installation or uninstallation of the Virtualization Engine instead of perf
wizard and how to use it, see Manual post uninstallation cleanup after
What's new as of December 30, 2005
Applying the latest fixes
It is important that you apply fixes as they become available. This top
changed.
Obtaining fixes - To access the Virtualization Engine fixes, s
Web site.

Figure 1: Information Center news

Figure 3-1 on page 63, Figure 3-2 on page 63, Figure 3-3 on page 64, and Figure 3-3 on page 64 describe the operating systems that the Virtualization Engine components support at the time of the writing of this redbook.

To get the latest information about the supported operating systems, the Information Center offers a view either by Virtualization Engine component or by operating systems, as shown in Figure 2 on page 347.

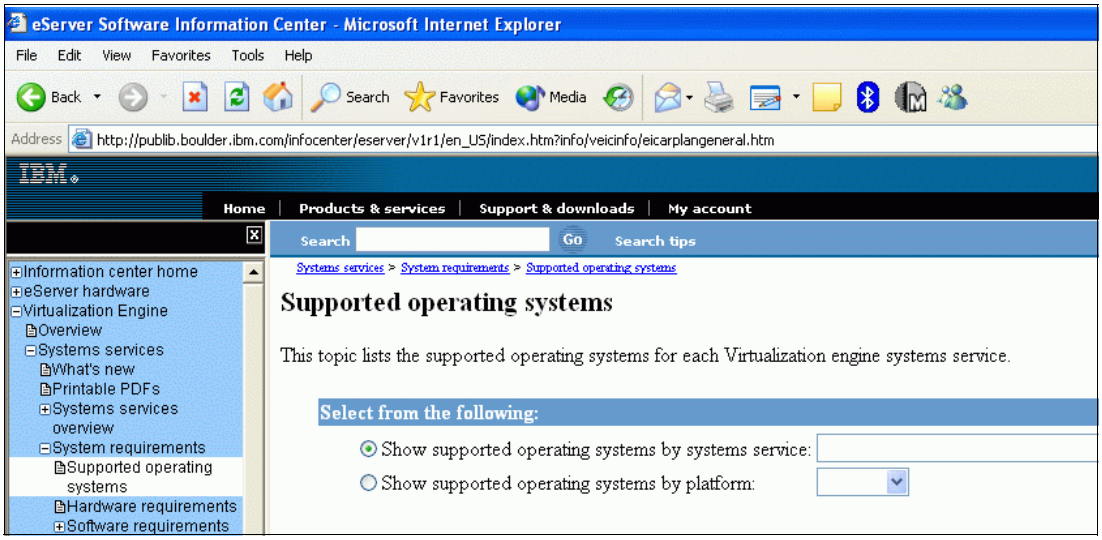
eServer Software Information Center - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address http://publib.boulder.ibm.com/infocenter/eserver/v1r1/en_US/index.htm?info/veicinfo/eicarplangeneral.htm
IBM
Home | Products & services | Support & downloads | My account
Search Go Search tips
Information center home
eServer hardware
Virtualization Engine
Overview
Systems services
What's new
Printable PDFs
Systems services overview
System requirements
Supported operating systems
Hardware requirements
Software requirements
Systems services > System requirements > Supported operating systems
Supported operating systems
This topic lists the supported operating systems for each Virtualization engine systems service.
Select from the following:
Show supported operating systems by systems service:
Show supported operating systems by platform:

Figure 2 Supported operating systems

Choose the right management server operating system

In the Virtualization Engine Version 1, AIX, OS/400, Linux, and Windows are the operating systems you can choose for most of the management servers components, offering a wide choice of possibilities.

- A first recommendation is to use one single operating system to set up all the management server components: that will facilitate their integration, and the skills required will be simplified. However in Virtualization Engine Version 1 it might happen that some of the components are not available, yet, on a specific platform and that you need to have a mix (for example as described in chapref).
- Another consideration might be the availability: if the management servers are down, you will not be able to manage the target environment. Today the Virtualization Engine Version 1 does not offer any backup facility.
- The security might also be an issue: be sure your management servers can be integrated at the security level you need/want into your environment.

A very important point is that all the operating systems provide the same functions for the managed servers: whatever the operating system you choose for the management server, the managed servers will be managed the same way (except for some very specific functions, for example from the IBM Director, which can offer particular functions, depending on the server type).

Print (and read) the documentation

There is already a lot of documentation covering the Virtualization Engine and its components and we provide the list in the , “Related publications” on page 389.

The Information Center provides as well a specific documentation you cannot get from another way and is provided as shown in Figure 3 on page 348.

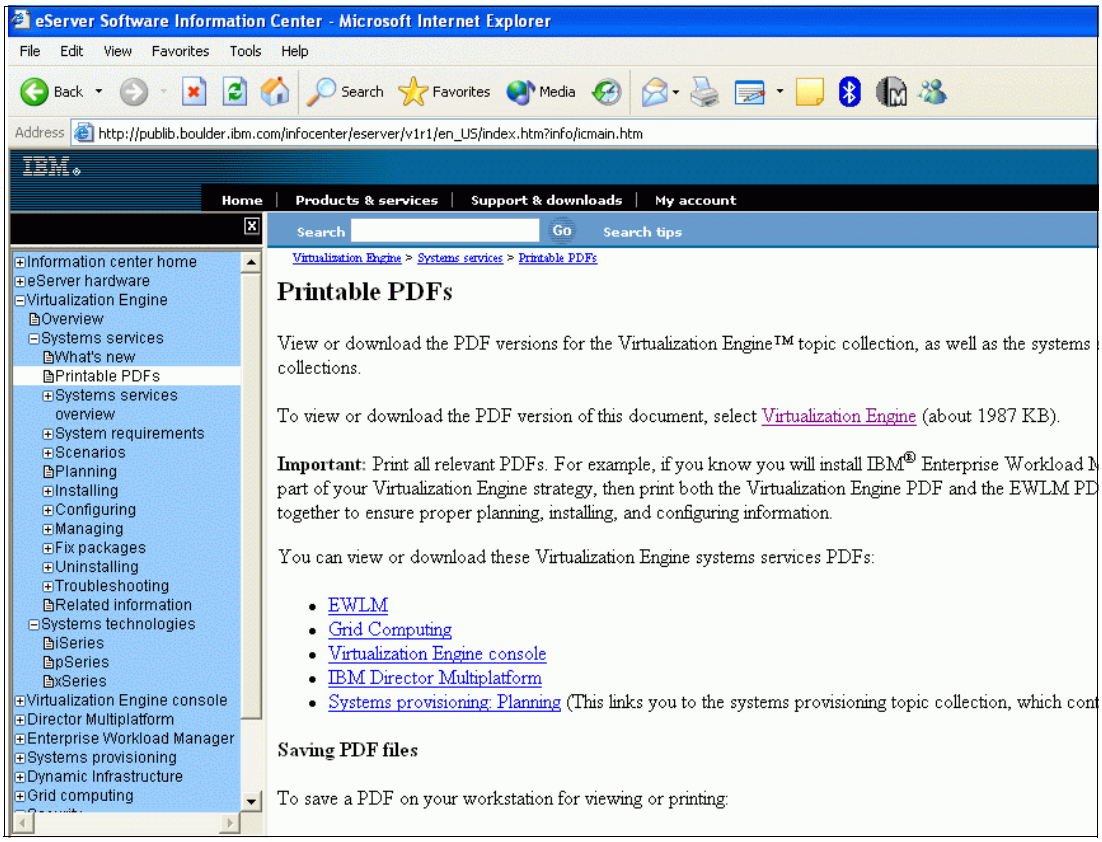
eServer Software Information Center - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address http://publib.boulder.ibm.com/infocenter/eserver/v1r1/en_US/index.htm?info/icmain.htm
IBM
Home Products & services Support & downloads My account
Search Go Search tips
Information center home
eServer hardware
Virtualization Engine
Overview
Systems services
What's new
Printable PDFs
Systems services overview
System requirements
Scenarios
Planning
Installing
Configuring
Managing
Fix packages
Uninstalling
Troubleshooting
Related information
Systems technologies
iSeries
pSeries
xSeries
Virtualization Engine console
Director Multiplatform
Enterprise Workload Manager
Systems provisioning
Dynamic Infrastructure
Grid computing
Virtualization Engine > Systems services > Printable PDFs
Printable PDFs
View or download the PDF versions for the Virtualization Engine™ topic collection, as well as the systems
collections.
To view or download the PDF version of this document, select Virtualization Engine (about 1987 KB).
Important: Print all relevant PDFs. For example, if you know you will install IBM® Enterprise Workload
part of your Virtualization Engine strategy, then print both the Virtualization Engine PDF and the EWLM PD
together to ensure proper planning, installing, and configuring information.
You can view or download these Virtualization Engine systems services PDFs:
EWLM
Grid Computing
Virtualization Engine console
IBM Director Multiplatform
Systems provisioning: Planning (This links you to the systems provisioning topic collection, which cont
Saving PDF files
To save a PDF on your workstation for viewing or printing:

Figure 3 Books available from the Information Center

We encourage you to make the list of the documentation that you need, to print it, and to read it before you start the installation. This will help you to avoid many problems.

The Information Center offers a very interesting dynamic tool for the installation and the configuration; however we encourage you to use both, the Information Center and the printed documentation altogether.

Be aware that the Virtualization Engine Version 1 does not support any national language version (except the IBM Director and Tivoli Provisioning Manager).

Using the planning advisor

As described in Chapter 4, “The planning advisor tool” on page 67, the planning advisor needs to be completed before you start the installation and the configuration. This wizard helps you:

- Identify what systems services you want to install.
- Design your Virtualization Engine topology.
- Generate customized installation flow charts.
- Order the components.

The Planning Advisor is available at:

`http://publib.boulder.ibm.com/infocenter/eserver/v1r1/en_US/index.htm?info/icmain.htm`

We recommend you print and record your customized output before beginning the installation process.

Verifying the system requirements

Be aware of the latest software and hardware requirements. Review the following information:

- Software requirements have been already mentioned. Figure 3-1 on page 63, Figure 3-2 on page 63,Figure 3-3 on page 64 and Figure 3-3 on page 64 express the requirements as today. We recommend that you look at the latest requirements as more and more operating systems become available.

  Fixes must be applied as well as stressed later on.

- Hardware requirements depend on many factors from your environment: the number of managed servers, the services you want to implement, the number of transactions you will monitor, the number of pools you want to create and manage, the number of applications you want to manage, the number of policies you want to implement, the number of workflow you will set up, and so forth.

  These requirements are improved by the experience and by the documentation of these experiences.

  As a starting point, we recommend the following (estimated size for moderate policy definitions supporting up to 50 managed servers):

  - For EWLM: 2 GB of disk space, whatever the platform, and a processor capacity in a range of 2 GHz for xSeries, 1.5 GHz for pSeries and 500 CPW for iSeries.
  - For Tivoli Provisioning Manager: 20 GB of disk space, whatever the platform, and a processor capacity in the range of 2.8 GHz for xSeries,1 GH for pSeries, and 450 CPW for iSeries.
  - For the Virtualization Engine console: 1 GB maximums of disk space, whatever the platform, and a processor capacity in the range of 1.5 GHz for xSeries, 1 GHz for pSeries, and 500 CPW for iSeries; 1 GB of main memory (expect 2GB on iSeries) will be needed for a good response time. Temporary disk space needs to be available as well: count at least 1 GB.
  - For the IBM Director Server, expect to use less than 1 GB of disk space maximum whatever the platform, and processor capacity around or less 1 GHz whatever the platform seems acceptable.
  - For the IBM Director Console, 500 MB of disk space and a processor capacity of 300 MHz seem affordable.

Tip: We recommend that you have a system that is of the right size before trying to install the IBM Virtualization Engine. Many problems occur when the management servers are too small. The Information Center is the right place to check for the latest best practices recommendations.

Performance consideration

Before installing you might want to validate the performance capacity of your configuration. The document *IBM Virtualization Engine Suite Performance Reference, Virtualization Engine Suite Version 1.1* gives you information to help to understand the performance and tuning factors of the Virtualization Engine components. This document is available at:

`http://publib.boulder.ibm.com/infocenter/eserver/v1r1/en_US/info/sc415018.pdf`

The intent of this document is to help provide guidance in terms of performance, capacity planning information, and tips to obtain best performance when using various components of the Virtualization Engine. This document includes performance information about the Virtualization Engine console, Enterprise, Workload Manager (EWLM), IBM Director Multiplatform, and associated products.

Complete the worksheets

Many components, everyone one with specific requirements need to be installed. To keep track of the server connection, security settings, bind information, ISC configuration, and so forth parameters, we recommend that you complete the worksheets in the planning advisor. Table 6-2 on page 121 is an example of such a worksheet.

We recommend you take the time to define these parameters before the installation process and to keep the worksheets up to date in the installation documentation workbook. Otherwise, if you forget a parameter (a password for example), you might need to reinstall the entire component.

Define your network requirements

When you install the Virtualization Engine, all the managed servers need to be connected to the management server and exchange some information.

Particular attention needs to be given to the network configuration parameters to avoid too many connection addresses and to avoid mixing production data with management systems data. For example, you might want to build different networks: one for the production data, and another one for the systems management data; this type of configuration can require additional network cards.

Identify product coexistence considerations

The following installation processes need specific consideration:

- The IBM Virtualization Engine Version 1 requires the installation of WebSphere Application Server V5.1. If you already have WebSphere Application Server V5.1 installed in your environment, you can use it. However, you need to be aware that the Virtualization Engine installation process will overwrite some of the installed parameters.
- The Virtualization Engine console requires, for a short period of time at the installation process, the port numbers of WebSphere Application Server V5.1. Thus, installing the Virtualization Engine console can cause any application using the WebSphere Application Server default ports to end.

Choose the code support

You can install the Virtualization Engine, either from physical media or from downloaded files:

- Physical media

  The installation process consists of two wizards: a Media Transfer wizard and an Installation wizard. The two wizards appear sequentially. The first part of the wizard copies the installation images to your system. Verify that your CDs and the operating system you are copying the images upon correspond to one another. The CD set you receive is operating system specific.

- Downloaded files

  You can download the Virtualization Engine files from Passport Advantage to a central location. As a result, these files can easily be updated and distributed to other systems from a central point. This is what we used in this project.

  The downloaded files need to be untarred or unzipped totally before you can use the Virtualization Engine Installation wizard which guides you through the process of installing your Virtualization Engine systems services.

In both cases, as you begin the wizard, we recommend that you have your completed planning worksheets and checklists available to help you complete the wizard pages correctly.

Apply the fixes

We recommend that you install the latest Virtualization Engine fix packs from the following Web site:

`http://techsupport.services.ibm.com/server/VirtualizationEngine`

From the site, you can download the appropriate fix packs for the Virtualization Engine common runtime and your Virtualization Engine systems services. The fixes can be downloaded by date, by operating system, by group, or by component as shown in Figure 4.

You install a fix pack by using a fix pack installation wizard, which is provided as an executable file within the archived fix pack that you download. The fix pack installation process is similar to the installation process for the Virtualization Engine itself.

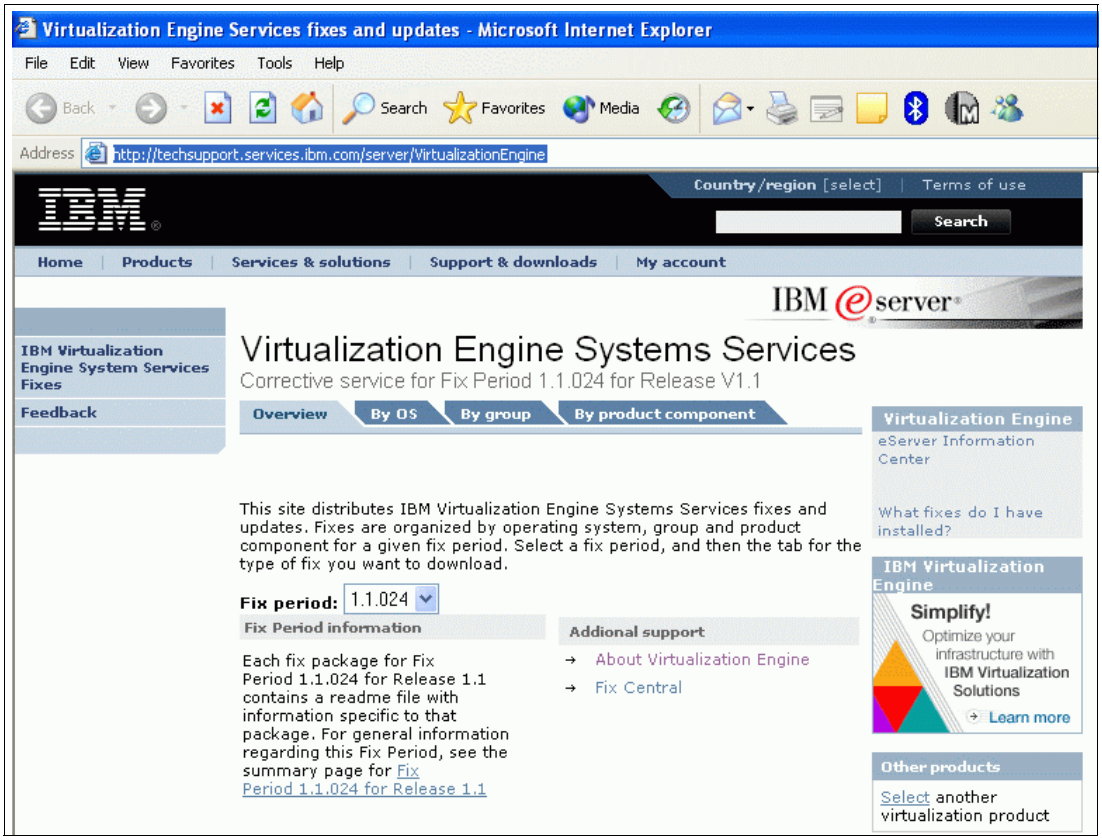
Virtualization Engine Services fixes and updates - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address http://techsupport.services.ibm.com/server/VirtualizationEngine
IBM
Country/region [select] Terms of use
Search
Home Products Services & solutions Support & downloads My account
IBM eserver
IBM Virtualization Engine System Services Fixes
Feedback
Virtualization Engine Systems Services
Corrective service for Fix Period 1.1.024 for Release V1.1
Overview By OS By group By product component
This site distributes IBM Virtualization Engine Systems Services fixes and updates. Fixes are organized by operating system, group and product component for a given fix period. Select a fix period, and then the tab for the type of fix you want to download.
Fix period: 1.1.024
Fix Period information
Each fix package for Fix Period 1.1.024 for Release 1.1 contains a readme file with information specific to that package. For general information regarding this Fix Period, see the summary page for Fix Period 1.1.024 for Release 1.1
Addional support
About Virtualization Engine
Fix Central
Virtualization Engine
eServer Information Center
What fixes do I have installed?
IBM Virtualization Engine
Simplify!
Optimize your infrastructure with IBM Virtualization Solutions
Learn more
Other products
Select another virtualization product

Figure 4 Virtualization Engine Systems corrective services

Verify the installation by component

We recommend that you verify the installation has been completed successfully at the end of each component installation. Should you have a problem, using this approach will make it easier to find the reason and to correct it quickly.

A

Planning for security

Note: In the ITSO project, we did not study the security possibilities when using the IBM Virtualization Engine. This appendix is the result of a project conducted in IBM France. The security aspects can be important for most of the installation. This is why we joined the results of this specific study.

In this appendix, we describe how to plan the global systems security and to configure it when using the following products:

- The IBM Virtualization Engine console on AIX 5.2
- The IBM Director Server on Windows 2000 Server sp4
- The IBM Director Agent on AIX 5.2, Windows 2000 Server sp4 and os400 V5R3
- The Enterprise Workload Manager Domain Manager and Managed servers on Windows 2000 Server sp4 and AIX 5.2

A.1 Planning

In this section, we describe the stakes of the Virtualization Engine security features. The most important point in a good security policy is the homogeneity. There are several layers and different domains to be secured in an information system and we can consider that the global security level is as high as the weakest link. That is why the Virtualization Engine security policy should be included in an already implemented global logic.

A.1.1 General consideration

A global security policy is composed of several domains:

- The servers' accesses:
  - Net access (DMZ, firewalls)
  - Login access
  - Services access (application accounts)
  - Files access (chmod, chroot)

© Copyright IBM Corp. 2005. All rights reserved.

- The data flows:
  - Secured application authentications (SSL certificates)
  - Transfer encryption (DES, 3DES…)
- Information protection:
  - Servers' login and password integrity (avoiding telnet connections for example)
  - Applications' services login/password integrity (limiting configuration files access)
  - Services' configuration integrity (limiting user accounts access)

Depending on your security policy, you can avoid some aspects such as encryption or secured application authentication but you have to keep in mind the global security framework. For example, a strong server access policy without any information protection is not part of the best practices because anyone who has an account (even with low privileges) can access the configuration files, and is able to find the administrators logins and passwords (LDAP configuration or IBM Director's administrator access).

The Virtualization Engine can be adapted to most of the security practices encountered in professional information systems. The suite supports Firewalls and Secured Networks, accounts management, Active Directory administration, SSL certificates, HTTPS consoles and Encrypted Data Transfers between all the elements.

A.1.2 Access and authentication

The first security step is to limit the accesses to the different servers and services.

You need to determine who is going to be allowed to log in to the servers and what are the privileges to be given, depending on each role. The different passwords should respect a special policy for example, by avoiding the dictionary words, by using upper and lower cases, and non-alphanumerical characters.

The Virtualization Engine supports many features to authenticate users, applications, management servers or managed servers. Most of the time, consoles, agents or servers' user management is based on the operating systems user accounts except the Virtualization Engine console which requires an LDAP authentication.

Between the different application services, an authentication can be implemented with SSL certificates or Lightweight Third Party Authentication keys (this is the case for example, between the IBM Director bridge and the Virtualization Engine console). It allows an integrity control of all elements of the architecture.

Figure A-1 describes the IBM Director bridge LTPA authentication process through the Resource Adapter for any user connection.

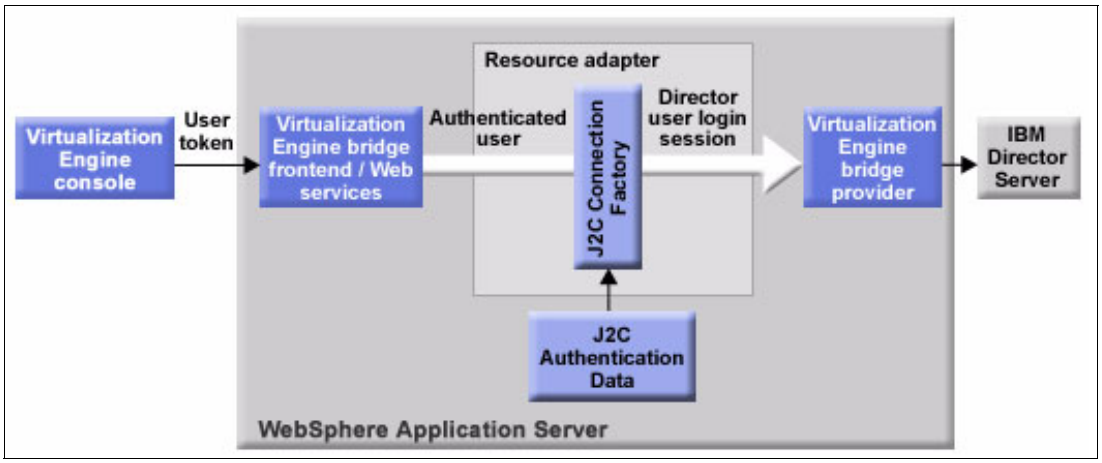

Figure A-1 The IBM Director bridge LTPA authentication process through the Resource Adapter

Finally, the infrastructure integrity can be guaranteed by administrating a network domain controlled with an Active Directory server.

A.1.3 Data encryption

Some data transfers can be considered as critical in the Virtualization Engine platform. Between each managed servers (agents) and management servers, the data flows contain information about the different systems:

- Logins and passwords
- Files
- Resources available
- Alerts
- Jobs running
- Performance statistics

Depending on the criticality of the managed servers, these data exchanges should be encrypted because any piece of information might be a global integrity threat during an intrusion attempt. If you plan to secure the servers' accesses you have to encrypt the communications between the IBM Director Servers and their Agents, and also between the Director Console and the management server.

It is also recommended that you implement SSL connection's between the EWLM Domain Manager and the Managed Servers if you are planning a higher protection against Deny of Service attacks for example.

A.1.4 Files and application servers access (data protection)

You might keep in mind that a good security policy continues with a data protection, insuring the integrity of your configuration. If anybody is allowed to edit the configuration files on the management servers, the Virtualization Engine R1 is vulnerable.

The IBM Director Server, for example, is as critical as the managed servers. Because of its functionality, capabilities and administrator rights on the agents, the Director Administrator user is very powerful. That is why it is very important to limit accesses to the management servers and the Bridges to the Virtualization Engine Consoles configuration files.

It's highly recommended to restrict the access to the Director Bridge configuration files containing LDAP server access and Director Administrator account.

The good attitude toward this topic is to limit access related to the real needs of each user role.

A.1.5 Virtualization Engine secured networks support

Each Virtualization Engine element supports secured network integration. Figure A-2 depicts a configuration example of a simple EWLM architecture.

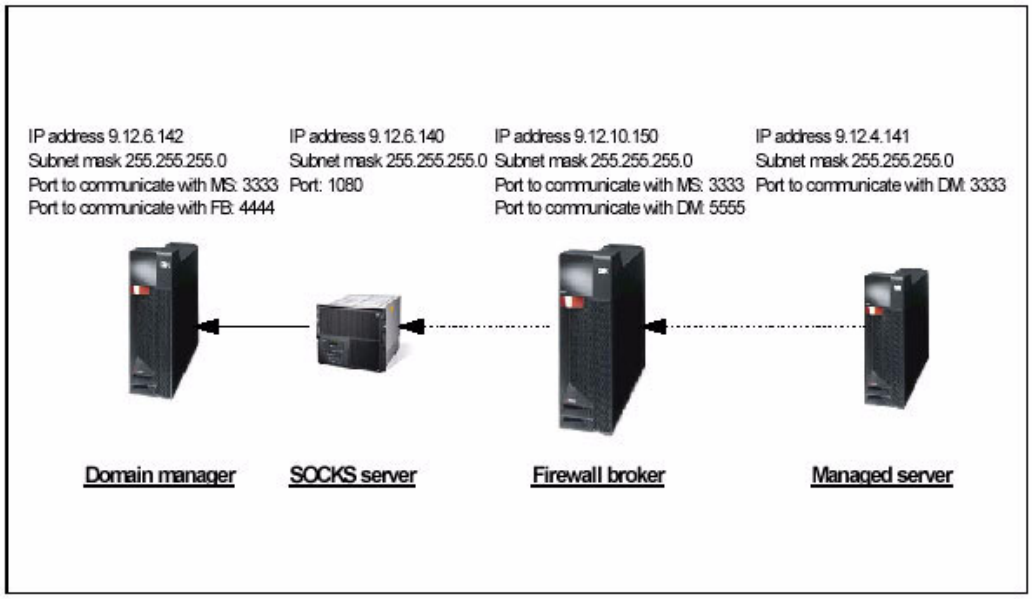

Figure A-2 A simple EWLM architecture.

EWLM supports any firewall or SOCKS server architecture through the Firewall Broker, a server controlling the communication ports use between the Domain Managers and all the Managed Servers.

The Virtualization Engine deployment impact on the firewall and SOCKS server configurations depends on the architecture. For example, in our EWLM configuration, as you can see on the previous figure, the SOCKS server has to be configured to accept the connections coming from the Firewall Broker to which all the Managed Servers are connected. This is due to the fact that there is no connection initialized by the Domain Manager itself.

IBM Director also supports this kind of DMZ configuration through the WBEM protocol.

A.2 Configuring

In this part, we describe the different ways to configure the security on each product, mainly command lines, files manipulations, configuration menu, or console menu. For more details, see *IBM Enterprise Workload Manager*, SG24-6350.

A.2.1 Virtualization Engine console

The Virtualization Engine console R1 is a view of the Director Server console configuration and a portal centralizing the accesses to all the other consoles. The user authentication is based on the Lightweight Directory Access Protocol server and the bridges authentication for the management sources is based on a Lightweight Third Party Authentication key.

Console to Management Source Bridges

During the installation of the Virtualization Engine Console, a LTPA key is generated. This file has to be used for any bridge installation. Another method is to get this file at the first bridge installation and import it in the Virtualization Engine Console or other bridge installation. Figure A-3 describes the LTPA authentication process.

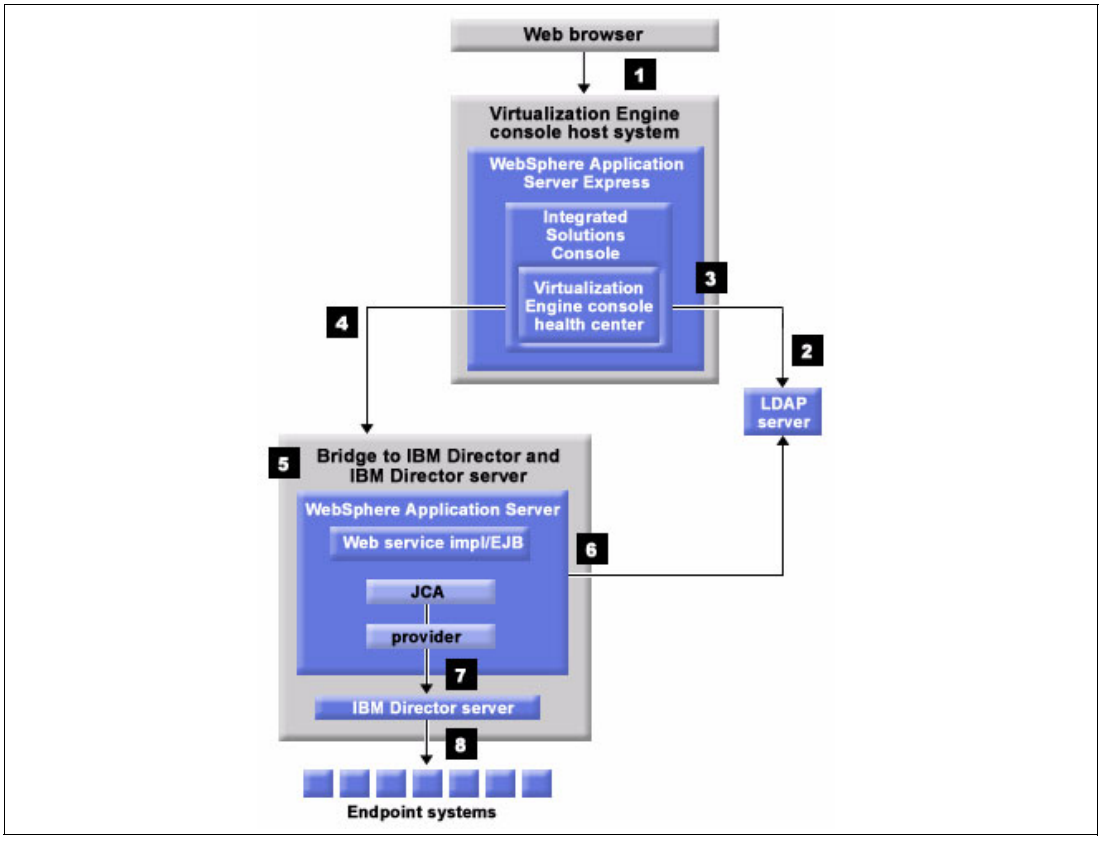

Figure A-3 The LTPA authentication process

In a typical usage scenario in which a system administrator performs management tasks on an endpoint, the course of events is as follows:

1. The authorized user connects to the ISC via a browser and is prompted to enter a login name and password.

2. The ISC is configured to authenticate the user against an LDAP user registry. If the authentication is successful, an LTPA token is created for the user and stored as part of the user session. The user has access to the Virtualization Engine console application.
3. The LTPA token contains the user identification and an expiration time. It is passed on to management sources as a credential, for example, all management sources must be part of the same LTPA realm. LTPA tokens will be configured with Single Sign On. This means that a set of keys is created on the Virtualization Engine console with which the LTPA token is encrypted. This set of keys must be propagated to all management source machines.
4. When making a call to the management source, the LTPA token is sent with the request. WebSphere Application Server will decrypt the LTPA token using the SSO keys, and verify the user's identity against LDAP. This means that both the Virtualization Engine console and the management source need to have a common LDAP configuration.
5. The WebSphere Application Server Container will authorize the user to the IBM Director Web services. Because these Web services have been implemented as Enterprise Java Beans (EJB™) under a security role, Admin, which is part of the EJB deployment descriptor, the default settings will be configured to allow all authenticated users to WebSphere Application Server to invoke operations on the EJB.
6. The EJB will invoke the Director JCA to get an authenticated Director Session. The WebSphere Application Server container will provide the mapped/pre-configured Director userid/password for the IBM Director Management server to the DirectorJCA. The DirectorJCA will use these credentials to acquire an authenticated Director session.
7. The rest of the components that make up the EJB implementation (for example, Provider code) will use this authenticated session to connect and receive information from the IBM Director Management server, including performing tasks.
8. Normal IBM Director server-to-agent and resource communication takes place (if necessary) to perform the requests by the Web services and EJB's on the management source. The results are passed to the provider module which translates the data into an XML response for the client. The Web services return the XML response document to the client. The GUI is updated with the new information.

HTTPS

SSL certificates are not supported in the first Virtualization Engine version. So, the Virtualization Engine Console is only available in http mode and the LDAP requests have to be unsecured.

A.2.2 IBM Director

As already discussed, IBM Director is a powerful administration server and needs to be well secured in order to avoid threating all the managed servers. IBM Director offers the following benefits:

- The console communications can be encrypted to protect remote sessions and authentications.
- The agents can be secured by login accounts asked to any IBM Director server. This prevents the system to be managed by another server implemented in a non-trusted remote area.
- The communications between the server and the agents (or any other Director servers) can be encrypted to protect any type of data flows. By default, the products support DES and 3DES protocols.

IBM Director can be vulnerable due to the bridge configuration files containing the login and passwords for the Virtualization Engine Console access. So, it is strongly recommended to install the bridge on a secured server with restricted access.

Console to Server

This section explains how to configure SSL communications between a director server and a console on Windows.

1. On the Console machine:
   a. Go to the DIRECTOR_HOME\data directory.
   b. Edit the TWGConsole.prop file.
   c. Modify, add or uncomment the following line:

      ```
      twg.gateway.link.1=com.tivoli.twg.libs.TWGSSLLink
      ```

2. On the Server:
   a. Go to the DIRECTOR_HOME\data directory.
   b. Edit the TWGServer.prop file.
   c. Modify, add or uncomment the following line:

      ```
      twg.gateway.link.1=com.tivoli.twg.libs.TWGSSLLink
      ```

3. Restart all the server services and the console.

Server to Agent

When you are installing an agent you can select Secured Mode or Data Encryption. Note that the encryption implies a secure mode. If an agent is already installed and managed by a server, you can open a console, right-click the managed system you want to secure, and click Secure System. The highest security level is the Data Encryption mode.

We are going to see how to configure it on all supported platforms:

- Windows:

  In the Control Panel Windows, click **Add/Remove Programs** then choose IBM Director Agent and click Change. Choose Modify and follow the procedure until you can tick Encrypt Data Transmission mode.

  The server then will have to be restarted.

- AIX:

  Note that the default agent installation directory is /opt/ibm/director.

  Run the following commands:

  ```
  cd /opt/ibm/director
  ./twgstop
  ./cfgsecurity
  ```

  Choose option 1.

  ```
  ./twgstart
  ```

- OS 400:

  Run the following commands:

  ```
  ENDTCPSVR *DIRECTOR
  QSH
  ```

  In the QSH prompt:

  ```
  /Qibm/ProdData/director/bin/cfgsecurity
  choose option 1
  STRTCPSVR *DIRECTOR
  ```

- LINUX:

  See the AIX method.

Note: For the WBEM support, refer to the Information Center at:

http://publib.boulder.ibm.com/infocenter/eserver/v1r1/en_US/index.htm?info/icmain.htm

The bridge to the Virtualization Engine console

The Virtualization Engine Console authenticates itself through the LDAP server on the bridge. Different users can be added following the procedure as described in the Information Center:

http://publib.boulder.ibm.com/infocenter/eserver/v1r1/en_US/index.htm?info/icmain.htm

In the Virtualization Engine console, select **Health Center** → **Setting up your Environment** → **Installing Bridges to Management Sources** → **Bridge to IBM Director** → **Post Installation Security Configuration**.

There is no data encryption between the Virtualization Engine console and the IBM Director Server (through the IBM Director bridge). If your operation system allows restricted access to some files, it is strongly recommended to only allow the bridge files access to administrator users and application services owner.

A.2.3 Enterprise Workload Manager

All the data flows are already encrypted between the Domain Manager and the Managed Servers but, in a strong security policy, SSL certificates can be implemented in order to authenticate any connection. By default, the Control Center can be accessed with an internet browser in HTTPS mode.

The Domain Manager and the Managed Servers support 2 SSL modes:

- Server authentication:

  A key and a certificate are generated on the Domain Manager, and then propagated on all the Managed Servers.

Figure A-4 describes the certificates distribution in SSL Server mode.

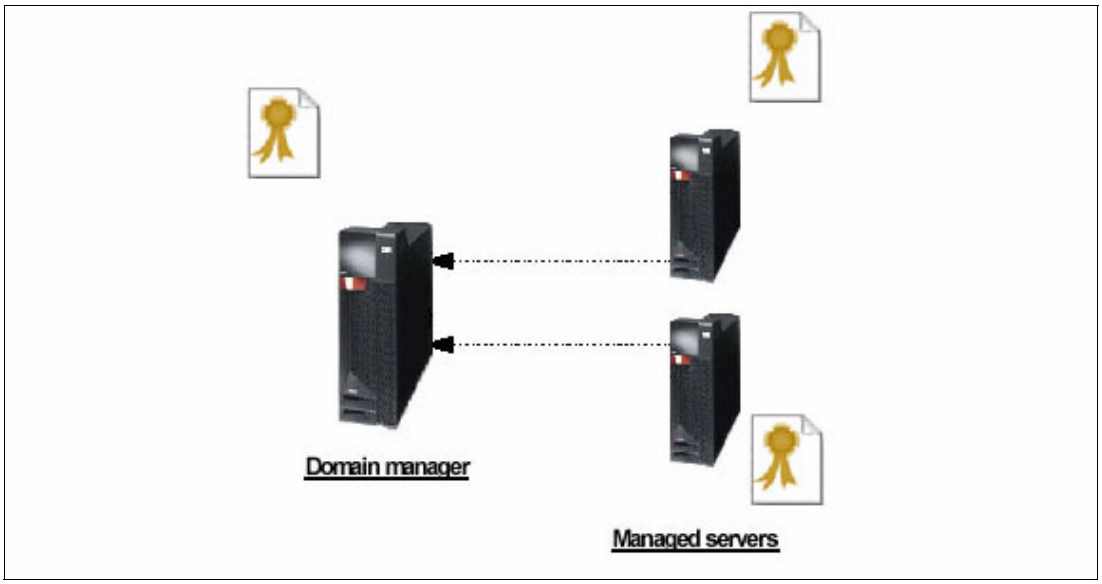

Figure A-4 The certificates distribution in SSL Server mode.Client/Server authentication

- Client/Server authentication:

  A key and a certificate are generated on the Domain Manager, and then propagated on all the Managed Servers. And each Managed Server generates its own key and certificate to be registered on the Domain Manager.

  Figure A-5 describes the certificates distribution in SSL Client/Server mode

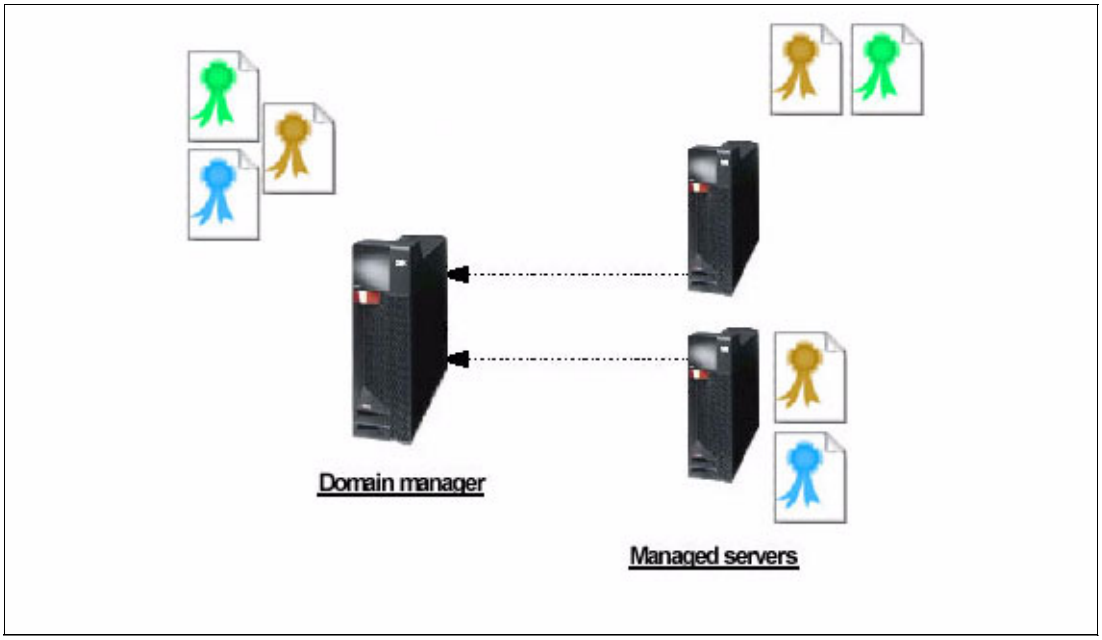

Figure A-5 the certificates distribution in SSL Client/Server mode

You can also combine the SSL protocol with a Firewall Broker configuration but between the Domain Manager and the Firewall Broker, a Client/Server SSL communication mode must be implemented, this is a requirement.

Between the Firewall Broker and the Managed Servers both mode are supported.

Figure A-6 shows the certificates distribution on the most complex configuration, a Client/Server SSL communication with a Firewall Broker:

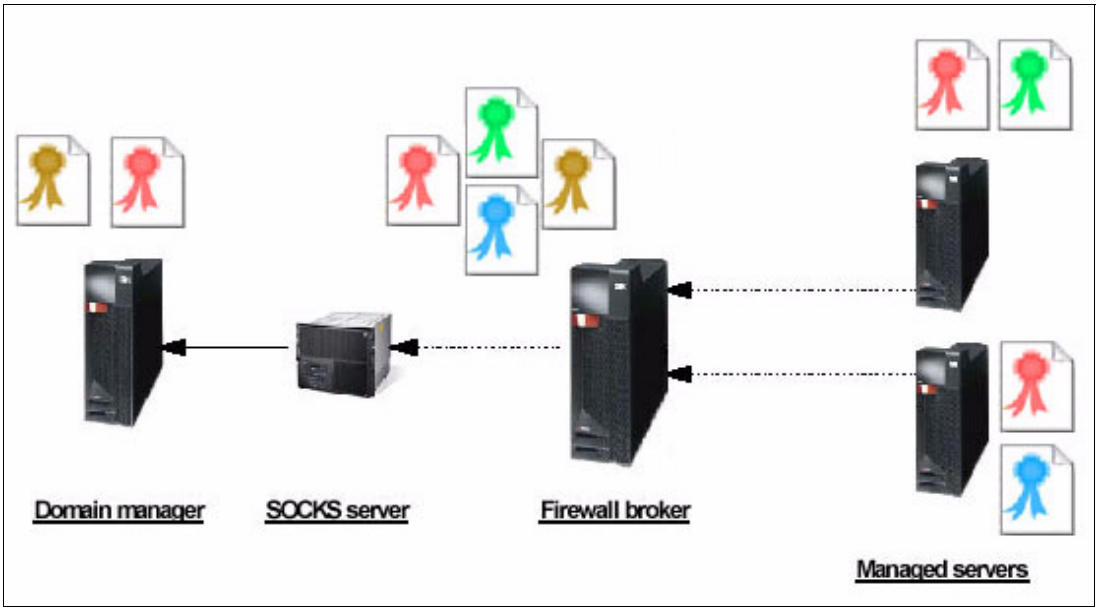

Figure A-6 A Client/Server SSL communication with a Firewall Broker

The Firewall Broker certificate is propagated on each server of this architecture. The Domain Manager and Managed Servers certificates are all registered on the Firewall Broker.

Now we are going to see how to configure both modes on each server.

Control Center to Domain Manager

The communication between the Control Center and the EWLM Domain Manager are secured through the Inter-Process Communication protocol.

Domain Manager to Managed Servers

In order to secure the servers authentication, SSL connections can be implemented. Depending on the security level you need, two modes are supported:

- The SSL Server mode, for a simple authentication. Each managed server recognizes the Domain Manager certificate.
- The SSL Client/Server mode, for client authentication. The domain server authenticates the registered managed servers and the managed servers authenticate the trusted domain server.

In this example, we describe how to configure these modes on an AIX Domain Manager, and Windows 2000 sp4 Managed Servers.

SSL Server configuration

To configure SSL Server, do the following:

1. At first, go in the EWLM Domain Manager /bin directory. (/opt/IBM/VE/EWLM/bin). A key has to be generated with the following command:

```
>keytool -genkey -alias dmAIX1 -keyalg RSA -keystore ks_dm -validity 365
```

The alias is linked to the key and the certificate. You will need to use it when you register the certificate on a target server. Complete the fields with your information and passwords.

2. Now export the key in a certificate used to identify the Domain Manager on each Managed Server using the following command:

   ```
   >keytool -export -alias dmAIX1 -keystore ks_dm -rfc -file ksDM.cert
   ```

3. Import the newly created certificate in the Domain Manager Keystore file with a different alias; this is necessary for internal communications: with the following command:

   ```
   >keytool -import -alias dmSELF -file ksDM.cert -keystore ks_dm
   ```

4. Copy the certificate on each Managed Server through FTP bin transfer for example in the following directory: EWLMMS_HOME_jvm\jre\bin

5. Import this certificate in a Keystore (created automatically if not existing), by running the following command in the same directory:

   ```
   >keytool -import -alias dmAIX1 -file ksDM.cert -keystore ks_ms
   ```

6. Enable SSL Server communication mode after stopping the Domain Manager:

   The default EWLM directory is /opt/IBM/VE/EWLM/bin and change the 111111 string with the code you entered in the certificate configuration:

   ```
   >./changeDM.sh /EWLM_Working_Dir/ -auth ServerSSL -sslks  /ks_dm -sslpw 111111
   ```

7. Verification:

   ```
   >./displayDM.sh /opt/EWLMDM/
   ```

   Notice the following line in the results:

   ```
   AuthorityLevel /auth (ServerSSL)
   ```

   Figure A-7 is the results of the **./displayDM** commands.

```
Processing displayDM request.  Please be patient as this may take a while...
WLMConfig - configurable property settings:
   ViaProxyPort/vp(null)
   TracePlugin/tlog(Off)
   :
   :
   SSLKeystore/sslks(/home/ibmewlm/ks_dm)
   ComponentTrace/ct(250)
   DomainManagerPort/mp(3333)
   :
   :
   SSLKeystorePassword/sslpw(password suppressed)
   DomainManagerAddress/ma(9.12.4.142)
   AuthorityLevel/auth(ServerSSL)
```

Figure A-7 Example of the displayDM command

   – Enable SSL Server communication mode after stopping the Managed Server:

     ```
     >EWLMMS_HOME\bin\changeMS.bat EWLM_Working_Dir -auth ServerSSL -sslks ks_ms
     -sslpw 111111
     ```

   – Verification: >EWLMMS_HOME\displayMS.bat EWLM_Working_Dir

     Notice the following line in the results:

     ```
     AuthorityLevel /auth (ServerSSL)
     ```

SSL Client/Server configuration

To configure SSL Client/Server, do the following:

1. To configure the Client/Server SSL mode you must have configured first the SSL Server mode as described in the previous part. You do not have to enable the SSL Server mode (the 4 last steps) if you are planning to enable the SSL Client/Server mode.
2. Then generate a key on each Managed Server and export it in a certificate:

   ```
   >keytool -genkey -alias ms1 -keyalg RSA -keystore ks_ms -validity 365
   >keytool -export -alias ms1 -keystore ks_ms -rfc -file ks_ms1.cert
   ```

   Each Managed Server is identified with a different alias in the keys and the certificates, for example: ms1, ms2…
3. Import each certificate with the corresponding aliases in the Domain Manager Keystore:

   ```
   >keytool -import -alias ms1 -file /home/EWLM/keystores/ks_ms1.cert -keystore ks_dm
   ```
4. Enable the Client/Server mode after stopping the Domain Manager:

   ```
   >EWLM_HOME/bin/changeDM.sh EWLM_Working_Dir -auth ClientServerSSL -sslks ks_dm
   -sslpw 111111
   ```
5. Verification:
   - >EWLM_HOME/bin/displayDM.sh EWLM_Working_Dir
   - Notice the following line in the results:

     ```
     AuthorityLevel /auth (ClientServerSSL)
     ```
6. Enable SSL Server communication mode after stopping the Managed Server:

   ```
   >EWLMMS_HOME\bin\changeMS.bat EWLM_Working_Dir -auth ClientServerSSL -sslks ks_ms
   -sslpw 111111
   ```
7. Verification:
   - >EWLMMS_HOME\displayMS.bat EWLM_Working_Dir
   - Notice the following line in the results:

     ```
     AuthorityLevel /auth (ClientServerSSL)
     ```

     Figure A-8 is the result of the **./displayMS** commands

```
Processing changeMS request.  Please be patient as this may take a while...

PROCESSING COMPLETE

root@ewlm4:/opt/IBM/VE/EWLMMS/bin>./displayMS.sh /opt/ewlmMS
WLMConfig - configurable property settings:
   ViaProxyPort/vp(null)
   TracePlugin/tlog(Off)
   InterBrokerPort/dp(null)
   :
   :
   SSLKeystorePassword/sslpw(password suppressed)
   DomainManagerAddress/ma(9.12.4.142)
   AuthorityLevel/auth(ClientServerSSL)
```

Figure A-8 Example of the displayMS command

Firewall Broker configuration command lines differences:

All the steps or command lines are the same for the Firewall Broker configuration, except the SSL activation command. Instead of using the **EWLM_HOME/bin/changeMS** or **changeDM**, you have to use the **changeFB** command with the same parameters.

B

VMware installation

In this appendix, we discuss the installation of a VMware ESX Server. The following information is based on the original VMware installation guide that is provided by the VMware product. For further information, go to the corresponding original documentation.

In the sections in this appendix, we install the following VMware products:

- VMware ESX Server (on which all virtual machines in our setup run).
- VMware VirtualCenter Server (the managing instance of all VMware servers in a management domain, running as an service on a server).
- VMware VirtualCenter Client on an existing IBM Director Console server (the VMware VirtualCenter Client is the user interface to the VMware VirtualCenter service).

B.1 Setting up VMware ESX Server

If you plan to use the VMM tools within a VMware environment the only supported VMware software product in managing guest operating systems is the VMware ESX server. The ESX Server admittedly has the most restrictions in respect of hardware requirements. Before planning the installation it is recommended to look at the VMware internet site to get the list of approved and certified hardware components. For example the server has to have either internal SCSI disks or attached storage such as NAS or SAN. There is no support for Parallel-ATA or Serial-ATA.

To get an idea about a sufficient hardware setup for the VMware ESX Server environment: We used in our scenario a server with a 2-way Pentium processor on board, 1 GigaByte memory and SCSI disk RAID 0 with 18 GigaByte space. Only one Ethernet card was used to provide the connection to the console of the VMware ESX Server and all the network traffic for the virtual machines (this setup is not recommended but it works fine in our system).

In the following section the installation of the VMware related hardware and software components will be described in more detail.

© Copyright IBM Corp. 2005. All rights reserved.

B.1.1 Planning VMware ESX Server installation

For the installation of a VMware ESX Server you need only the bare metal. All previous software installations and data on that server will be erased in the partitioning step if you chose the default. If you not plan to donate the full disk for this installation, consult the VMware ESX Server documentation from VMware how to do it.

The following information is needed installing a VMware ESX Server. These parameters are required throughout the installation process:

- VMware ESX Server serial numbers that look similar to `12345-12345-12345-12345`
- Network definition for VMware ESX Server console:
  - Host Name: `msnesx1`
  - IP Address: `9.12.4.192`
  - Netmask: `255.255.254.0`
  - Gateway: `9.12.4.92`
  - Primary DNS: `9.12.6.7`
- The password for the root account on that VMware ESX Server.

We installed the VMware ESX Server with the corresponding installation CD. Be sure that you have this CD available to boot the server boot from this media.

B.1.2 Installing VMware ESX Server

Put the VMware ESX Server installation CD into the Compact Disk drive and boot the server.

The VMware ESX Server installation (we have used the graphical installer) guides you through different screens. Only the major steps that are important to define the parameters that configure the server instance are listed here. For example, one screen similar to the End User Licence Agreement screen was left out. For more information about installing VMware ESX Server, go to the software vendors home page and download the appropriate documentation. To move through the process, click **Next**.

These are the major configuration steps:

1. Installation Type screen: Choose **default**
2. VMware ESX Serial Numbers screen: Enter ESX Server licence key
3. Device allocation screen:
   a. Choose the default of 192 MB for **Reserved Memory**. This amount of memory is high enough to run up to 8 virtual machines on that server -- you could change this setting any time later through the management console of the VMware ESX Server if there is a demand to define a higher amount of memory to the server.
   b. If you have only one hard disk drive installed in your server than select **Shared with Service Console** to be able to store virtual machines images on the same disk on which the VMware ESX Server software is installed. If you have more than one disk unchecked this box for performance reasons to boost the disk trough put.
4. Disk Partitioning Setup screen: Choose **Automatic Partitioning** with the secondary option **Remove all partitions**.
5. Network Configuration screen: Choose **Static IP** and enter the well known network configurational parameters provided by your network administrator: Hostname, IP Address, Netmask, Gateway and Primary DNS.
6. Account Configuration screen: Set a password for the primary administrator account `root` on this VMware ESX Server: `password`.

The last configuration step has defined the server account information. Now the installation procedure for VMware ESX Server formats the disks and starts installing the Linux packages. The time needed for the unattended server installation takes between 5 to 15 minutes, depending on the speed of both the CPUs and the hard disks in the system.

On the Congratulation screen, the completeness of the installation process is displayed. Select **Next** and the server reboots.

Note: Every time the VMware ESX Server starts you have the chance to observe the progress and status of the boot sequence. If everything is starting as expected, the prompt of the steps initializing a daemon or a function is reported with `[OK]`. If there is an error, it might display `[FAILED]` or additional error messages might occur. After having installed the IBM Director Agent on a VMware ESX Serve, r one additional status line is displayed on the boot sequence: `Starting IBM Director Agent`. This indicates that the service has been installed successfully.

B.1.3 Verifying VMware ESX Server installation

After the successful installation of the VMware ESX Server you could launch the VMware Management Interface through a Web browser. The needed URL is displayed on the VMware ESX Server screen after the Linux kernel boot sequence of the server is completed. If this screen is not displayed you could switch to it by pressing the <Alt>-<F1> keyboard combination. The URL should now show up.

The URL looks similar to `http://msnesx1`, where `msnesx1` is the host name you have defined in the installation of that new VMware ESX Server instance. Be sure that the host name could be resolved through a DNS server query on your system. Depending on your setup you have to use the full qualified host name similar to `msnesx1.itso.ibm.com` instead of the host name itself such as `msnesx1`.

For a list of supported Internet browsers, consult the VMware support for VMware ESX Server.

Note: If you are using an Internet Explorer to access the VMware Management Interface be sure that the installed version is 6.0 or later. Otherwise, you get a runtime error of an unexpected quantifier. The only two options you could take is either to enter the debugging mode or to cancel the operation — both will not solve the browser issue.

Through the URL `http://msnesx1` you request access to the VMware ESX Server Management Interface. Log in to the system with the default user root and the corresponding password or with any other user you have defined through the installation process of the VMware ESX Server.

The VMware Management Interface provides you all functions you need to get status information and the ability to change system behavior through setting parameters. To structure the content logically the management interface offers you the choice of three flags:

Status Monitor: Get all information about the VMware ESX Server and the virtual machines running on that server

Memory: Get all system information about physical and reserved memory used and available to the system. Also a memory summary displayed by virtual machine running on that server

Options: Query and change all available parameters and system status information for this VMware ESX Server

Manage Files: Is an additional flag (in shape of an underlined URL link) on the top of the browser (beside the already mentioned three flags on the left hand side). This link opens the VMware File Manager. Now you have the access to the VMware ESX Server file system and you could browse through it and make property changes to the files on the system.

Hint: If you have installed the VMware ESX Server on a system with only one single disk you have to check if the system allows you to store also virtual machine images on the same disk. To verify this open the **Options** flag in the VMware Management Interface. Chose **Startup Profile...** and a new window opens. In the section **Hardware Profile** look at the **SCSI storage controller group** that the check box **Shared with Service Console** is ticked (see figure Figure B-1).

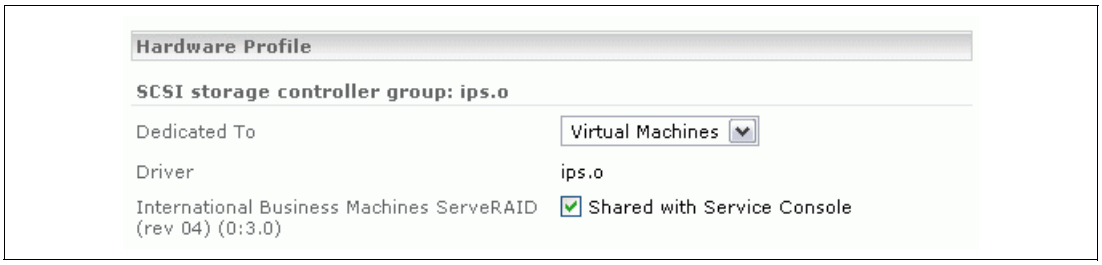

Figure B-1 SCSI storage controller group

B.2 Setting up VMware VirtualCenter

The purpose of VMware VirtualCenter is to have one central managing instance for all your VMware ESX and GSX Servers running in your environment.

For detailed information about the server hardware requirements for VirtualCenter Server we refer to the documentation provided by VMware.

In our setup we use the VMware VirtualCenter server to manage one VMware ESX Server instance, installed on `msnvmw1` as described in section B.1.2, “Installing VMware ESX Server” on page 368. The server running the VMware VirtualCenter server service is on `vewin22` (this is the Windows 2003 server running the EWLM that is described in Chapter 7, “Installing the Virtualization Engine console, IBM Director, TPM, and EWLM on Blade with Windows” on page 161.

B.2.1 Planning VMware VirtualCenter server installation

Due to the fact we are installing VMware VirtualCenter server on an already existing Windows 2003 server with Service Pack1 introduced in 7.4, “Installing Enterprise Workload Manager” on page 186 if you need detailed information.

The following information are needed installing a VMware VirtualCenter server. These parameters will be asked throughout the installation process:

- VMware VirtualCenter server database: If you plan to use a different database format for storing the VMware VirtualCenter information than the default (Microsoft Access) than you have to prepare your systems according to the VMware VirtualCenter preparation for databases (supported database SQL Server through an ODBC Connection, Oracle 8i or Oracle 9i).

- The VMware VirtualCenter server installer will ask for Customer Identification Information, be sure you have some at hand:

  User Name: `vmwadmin`
  Organization: `IBM`

- The VMware VirtualCenter Web service will run using an account that must be a member of the Administrators group on the server. For this purpose we suggest to create a new user with the following characteristics:

  User name: `vmwadmin`
  Member of: `Administrators`

- The VMware VirtualCenter Web Service will listen to a specific port. Scrutinize that this port is not already used by other programs or services on the server:

  Web Service port `8443`

B.2.2 Installing VMware VirtualCenter server

The VMware VirtualCenter server installation will guide you through different screens. Only the major steps that are important to define the parameters that configure the server instance are listed here. For example, one screen similar to the End User Licence Agreement screen was left out. For more information about installing VMware VirtualCenter server, go to the software vendors home page and download the appropriate documentation. To move through the process, click **Next**.

These are the major configuration steps:

1. Start the VMware VirtualCenter server installation program (on Windows systems an executable file in the VM ware VirtualCenter install directory). Depending on your media this could be on a CD or a subdirectory on your servers file system.
2. Customer Information screen: Enter the **User Name** `vmwadmin` and the **Organization** `IBM`.
3. Database Setup screen: Chose **Use Access database** if you allow VMware VirtualCenter server to manage his database information internally. For clarification: VMware VirtualCenter server only stores the information in a Microsoft Access compliant format; it does not use the Microsoft Access database product for doing so. The option **Use Access database** is only recommended if you are running in a small experimental environment. The manufacturer recommends to build production environments on one of the other database choices.
4. VirtualCenter Web Service screen: Enter the **User Name** `vmwadmin`, his **Password**, keep the default **Web Service Port** `8443` and select the **Use default VMware Digital Certificates** check box.
   This User Name plus Password plus Port data combination is the key for all Web Service related communications to this VMware VirtualCenter server instance.
5. Installation Completed screen: The installation is done you could exit the installation program by selecting **Finish**.

The time needed for copying of files, installing services and configurations tasks on the Windows operating system takes between one to two minutes, depending mostly on the speed of the hard disks in the system.

B.2.3 Verifying VMware VirtualCenter server installation

This section provides you some ways in checking if the installation of the VMware VirtualCenter server features succeeded successfully:

- From a program installation perspective: Select **Start** → **Control Panel** → **Add or Remove Programs**. You find one row that points to the installation of the VMware VirtualCenter. Chosing the link **Click here for support information** in that row and a new **Support Info** window opens and shows you the product information. The most useful information here is the version of VMware VirtualCenter (see Figure B-2).

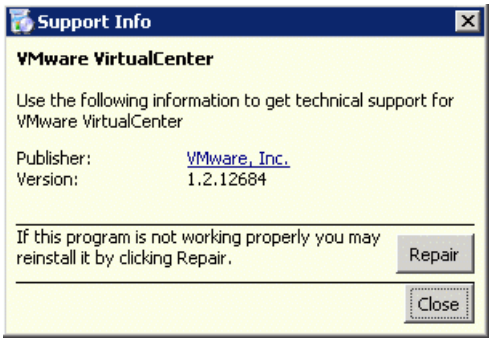
Support Info
VMware VirtualCenter
Use the following information to get technical support for VMware VirtualCenter
Publisher: VMware, Inc.
Version: 1.2.12684
If this program is not working properly you may reinstall it by clicking Repair.
Repair
Close

Figure B-2 VMware VirtualCenter Version 1.2.12684 is installed on this server

- From a windows system service perspective: Select **Start** → **Control Panel** → **Administrative Tools** → **Services** you start the program that monitors all installed services on that Windows system. The installation of the VMware VirtualCenter server has installed two important services:
  - Service name: vpxd is the VMware VirtualCenter Server that provides centralized management of VMware virtual machines. The startup type is Automatic and the service status is Started. On the Dependencies tab you could see that the VMware VirtualCenter Web Service is dependent on this vpxd service.
  - Service name: vmaService is the VMware VirtualCenter Web Service that provides a Web services interface for centralized management of VMware virtual machines. The startup type is Automatic and the service status is Started.
- From a VMware VirtualCenter Web Service perspective: The VMware VirtualCenter Installation guide provides a full chapter on that topic but we will introduce one quick way of verifying the successful installation:
  - Open a browser window on the server you have installed VMware VirtualCenter server. Type in the address field `https://localhost:8443/login.html` where `8443` is the VMware VirtualCenter Web Service port defined at installation time.
  - The VMware VirtualCenter Web Service will display a form in which the following values has to be typed in:

    User name: `vmwadmin`
    User password: `password`
    Web Service URL: `https://localhost:8443`
  - Click **Login** and in the browser window the top level information of that Web Service instance will be displayed in XML format (see Figure B-3 on page 373). The connection to the VMware VirtualCenter via Web Services was successfully established. The service is running.

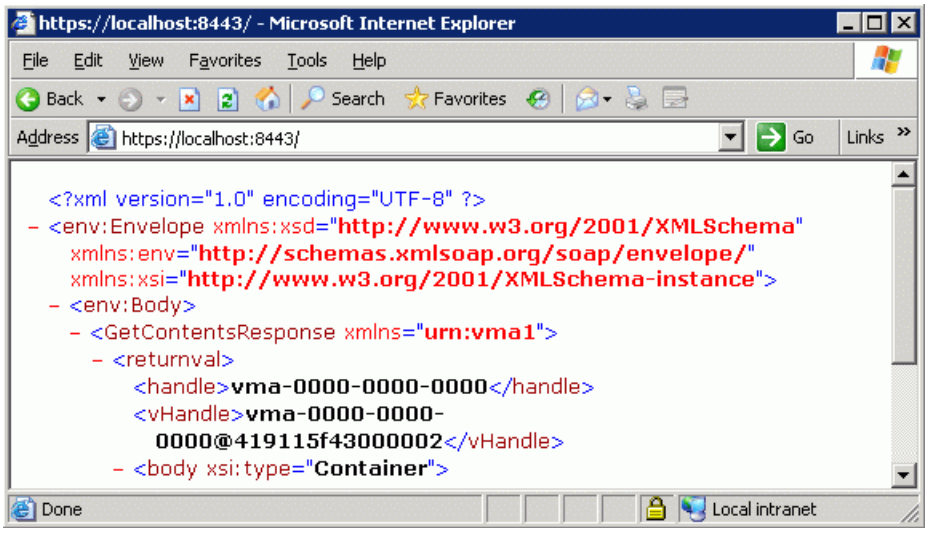

Figure B-3 Browser window with VMware VirtualCenter Top-Level Information in XML format

Note: If you are connected to the VMware VirtualCenter Web Service you could query the structure of that XML file by adding one identifier (see XML structure in you open browser) of type `<name>` to the URL. For example, if you would like to get the object `customProperty` displayed you have to add this name to the URL similar to `https://localhost:8443/customProperty`. The browser then shows the attributes available for `customProperty`.

At this point of the installation you have installed and verified that the VMware VirtualCenter server is installed, running and that the additional software package of Web Services are also installed and working properly. To get a direct access to the VMware VirtualCenter service you have to install in the next step VMware VirtualCenter Client on one system.

B.3 Setting up VMware VirtualCenter Client

The purpose of VMware VirtualCenter Client is to provide a user interface to the VMware VirtualCenter service running on a given server. For detailed information about the server hardware requirements for VirtualCenter Client we refer to the documentation provided by VMware.

In our setup we use the VMware VirtualCenter Client to manage one VMware ESX Server instance, installed on `msnvmw1` described in section B.1.2, “Installing VMware ESX Server” on page 368. The server running the VMware VirtualCenter server service is on `vewin22` (this is the Windows 2003 server running the EWLM.

The new component VMware VirtualCenter Client could be installed on any system that has a network connection to the VMware VirtualCenter service. In our environment we install all application of type console on the server `vewin21`. On this server also the IBM Director Console and the IBM Virtualization Engine console are installed.

B.3.1 Planning VMware VirtualCenter Client installation

The VMware VirtualCenter Client is a program that has from a prerequisite stand point of view no dependencies. It could be installed on every computer in the network.

The following information are needed installing a VMware VirtualCenter Client. These parameters are required throughout the installation process:

- The VMware VirtualCenter Client installer will ask for Customer Identification Information, be sure you have some at hand:

  User Name: `vmwadmin`
  Organization: `IBM`

Remark: The VMware VirtualCenter Client on a Windows 2000 system needs as a prerequisite the Microsoft .NET Framework version 1.1. If no .NET Framework is installed on the server the installer will launch the installer for the Framework. If and older version of .NET is already installed the installer will upgrade this version to 1.1.

B.3.2 Installing VMware VirtualCenter Client

The VMware VirtualCenter Client installation will guide you through different screens. Only the major steps that are important to define the parameters that configure the server instance are listed here. For example, one screen similar to the End User Licence Agreement screen was left out. For more information about installing VMware VirtualCenter Client, go to the software vendors home page and download the appropriate documentation. To move through the process, click **Next**.

These are the major configuration steps:

1. Start the VMware VirtualCenter Client installation program (on Windows systems an executable file in the VM ware VirtualCenter install directory). Depending on your media this could be on a CD or a subdirectory on your servers file system.
2. Customer Information screen: Enter the **User Name** `vmwadmin` and the **Organization** `IBM`.
3. Setup Type screen: Chose **Client** to install the Console
4. Installation Completed screen: The installation is done you could exit the installation program by clicking **Finish**.

The time needed for copying of files and configurations tasks on the Windows operating system takes between one to tow minutes, depending mostly on the speed of the hard disks in the system.

B.3.3 Verifying VMware VirtualCenter Client installation

The installation process creates a shortcut on the Windows desktop with the label VMware VirtualCenter. Launch the VMware VirtualCenter Client through this shortcut on the desktop. The first thing you have to do is the authentication against the VMware VirtualCenter server. The information needed are:

1. Server name in form of an IP address (9.12.4.179) or a logical name (vewin22) that could be resolved through the DNS server or the local hosts table.
2. The Username (vmwadmin) you have defined through the VMware VirtualCenter server installation process.
3. The Password (password) you have defined through the VMware VirtualCenter server installation process.

Starting the VMware VirtualCenter Client for the first time you have to enter licenses for all systems you will manage through your VMware VirtualCenter instance. Therefore the first dialog that opens is the ”VMware VirtualCenter Licensing” screen to enter the needed license keys:

1. VMware VirtualCenter Licensing screen: Click **Add Licenses** and a Windows operating system Open dialog opens.
2. Open screen: Navigate to the text file on the system disk that is holding the license keys for one or many components. Select **Open** and all valid licenses keys are imported and displayed in the window.
3. VMware VirtualCenter Licensing screen: Check if the license keys that you intended to be loaded into the system are there. Click **Done**.

Note: Only with a valid VMware VirtualCenter license key you are allowed to start the VMware VirtualCenter Client.

4. Now the VMware VirtualCenter Client window is open and ready for configuration.

At this point the VMware VirtualCenter Client is running, the Client has a set of valid VMware licenses. The next step is to build the physical and logical view of the different VMware servers in you managing domain. The following description for the definition of farms, hosts, virtual images will take place in the VMware VirtualCenter Client and is, therefore, the logical next step after installing license keys for VMware.

The first step is to define a farm. A farm is a logical construct to help structuring a server farm. For instance you could use the farm object to differentiate servers in you infrastructure by creating farm objects for Development, Test and Production as follows:

1. In the navigator on the left hand side, right-click the object **Server Farms** and select **New Farm.**
2. Below the object Server Farms in the navigator a new object of type Farm is created and you have to enter a name that identifies your farm: **`VE Farm`**.
3. The new farm object is created and you have access to that object through the navigator.

The second step is to define a host. A host is a server, for example a VMware ESX server, that is managed through VMware VirtualCenter. A host is a secondary object to a farm object. A host could not be added to a server farm as follows:

1. In the navigator on the left hand side right-click the one step before created object **VE Farm** and select **Add Host... to launch the VMware VirtualCenter Add Host Wizard.**
2. Wizard Welcome Screen screen: Click **Next**.
3. Specify Connection Settings screen: Enter parameters for the connection that point to a managed server and the authentication to get administrator access to the managed server.

Host Name:	`msnesx1` The host name could be a logical name or an IP address
Port:	`902`
User Name:	`root`
Password:	`password`

   In this setup you have to enter one of the administrator accounts on the VMware ESX Server configured in B.1.2, “Installing VMware ESX Server” on page 368. In our example we took the root account. Click **Next.**

Now the wizard is building up a network connection to the defined server and check out if the account parameter are valid for that server. If something is missing or wrong an error dialog will show up and tell what went wrong. If everything went well go to the next step:

4. Enable VMotion screen: In our setup we do not like to enable VMotion on this host because we do have only one physical hardware box with no partition options on that. Therefore it makes no sense to enable the mover from one box to another. Click **Next**.
5. Wizard Completion Screen screen: Click **Finish**.

A new host object msnvmw1.itso.ibm.com is added under your farm object VE Farm in the navigator. The name of this host object is the host name you have defined in the previous step for setting up a farm object. If there are already virtual machine images on that server, VMware VirtualCenter discovers those and creates a new group object Discovered VMs and put each VM as an new VM object under the Discovered VMs object.

For more information how to use VMware VirtualCenter Client, go to the VMware Web site and download the corresponding documentation.

C

Sample EWLM domain policy

This appendix provides a description of the EWLM domain policy that we used to set our goals for our Trade3 workload and monitor them.

In is a copy of the domain policy that we used for Trade3. We also performed some testing with the plants workload previously, so our domain policy includes service class definitions for both. Although the definitions were included for plants, this did not impact our scenarios as there was no work that entered the system for the plants workload.

Example: C-1 EWLM domain policy

```
<?xml version="1.0" encoding="UTF-8"?>

<DomainPolicy xmlns="http://www.ibm.com/eWLM/DomainPolicy">
  <Name>DC Trade3 and Plants Domain Policy</Name>
  <Description>DC Sample Policy for Redbook.</Description>
  <Version>1.00.00</Version>
  <UI>EWLM Editor</UI>

  <Workloads>

    <Workload>
      <Name>EWLM Workload</Name>
      <Description>Workload definition</Description>
    </Workload>

    <Workload>
      <Name>WL_shopping</Name>
      <Description>Shopping workload for Trade3 and Plants.</Description>
    </Workload>

    <Workload>
      <Name>WL_default</Name>
      <Description>Default workload for Trade3 and Plants.</Description>
    </Workload>

    <Workload>
      <Name>WL_Trade3_others</Name>
      <Description>Other workload for Trade3.</Description>
```

© Copyright IBM Corp. 2005. All rights reserved.

```
  </Workload>

  <Workload>
    <Name>WL_Plants_others</Name>
    <Description>Other workload for Plants.</Description>
  </Workload>

</Workloads>

<ServiceClasses>

  <ServiceClass>
    <Name>SC_Calc_batch</Name>
    <Description>Calc Batch service class for AIX</Description>
    <WorkloadName>EWLM Workload</WorkloadName>
  </ServiceClass>

  <ServiceClass>
    <Name>SC_InternetExplorer</Name>
    <Description>Internet Explorer service class for Windows 2003 Server.</Description>
    <WorkloadName>EWLM Workload</WorkloadName>
  </ServiceClass>

  <ServiceClass>
    <Name>SC_Plants_shopping</Name>
    <Description>Shopping service class for Plants.</Description>
    <WorkloadName>WL_shopping</WorkloadName>
  </ServiceClass>

  <ServiceClass>
    <Name>SC_Plants_default</Name>
    <Description>Default service class for Plants.</Description>
    <WorkloadName>WL_default</WorkloadName>
  </ServiceClass>

  <ServiceClass>
    <Name>SC_Plants_login</Name>
    <Description>Login service class for Plants.</Description>
    <WorkloadName>WL_Plants_others</WorkloadName>
  </ServiceClass>

  <ServiceClass>
    <Name>SC_Trade3_quotes</Name>
    <Description>Quotes service class for Trade3.</Description>
    <WorkloadName>WL_Trade3_others</WorkloadName>
  </ServiceClass>

  <ServiceClass>
    <Name>SC_Trade3_portfolio</Name>
    <Description>Portfolio service class for Trade3.</Description>
    <WorkloadName>WL_Trade3_others</WorkloadName>
  </ServiceClass>

  <ServiceClass>
    <Name>SC_Trade3_default</Name>
    <Description>Default service class for Trade3.</Description>
    <WorkloadName>WL_default</WorkloadName>
  </ServiceClass>

  <ServiceClass>
```

```
    <Name>SC_Trade3_profile</Name>
    <Description>Profile service class for Trade3.</Description>
    <WorkloadName>WL_Trade3_others</WorkloadName>
  </ServiceClass>

  <ServiceClass>
    <Name>SC_Trade3_account</Name>
    <Description>Account service class for Trade3.</Description>
    <WorkloadName>WL_Trade3_others</WorkloadName>
  </ServiceClass>

  <ServiceClass>
    <Name>SC_Trade3_shopping</Name>
    <Description>Shopping service class for Trade3.</Description>
    <WorkloadName>WL_shopping</WorkloadName>
  </ServiceClass>

  <ServiceClass>
    <Name>SC_Plants_general</Name>
    <Description>General service class for Plants.</Description>
    <WorkloadName>WL_default</WorkloadName>
  </ServiceClass>

  <ServiceClass>
    <Name>SC_Trade3_general</Name>
    <Description>General service class for Trade3.</Description>
    <WorkloadName>WL_default</WorkloadName>
  </ServiceClass>

  <ServiceClass>
    <Name>EWLM Service Class</Name>
    <Description>Service class definition</Description>
    <WorkloadName>EWLM Workload</WorkloadName>
  </ServiceClass>

</ServiceClasses>

<ServicePolicies>

  <ServicePolicy>
    <Name>DC Trade3 and Plants Service Policy</Name>
    <Description>Goal definitions for each service class for Redbook.</Description>
    <ServicePolicyGoals>
      <ServicePolicyGoal>
        <ServiceClassName>SC_Calc_batch</ServiceClassName>
        <Goal>
          <Velocity>
            <Importance>Medium</Importance>
            <Level>Slowest</Level>
          </Velocity>
        </Goal>
      </ServicePolicyGoal>
      <ServicePolicyGoal>
        <ServiceClassName>SC_InternetExplorer</ServiceClassName>
        <Goal>
          <Velocity>
            <Importance>Medium</Importance>
            <Level>Fast</Level>
          </Velocity>
        </Goal>
```

```
</ServicePolicyGoal>
<ServicePolicyGoal>
  <ServiceClassName>SC_Plants_shopping</ServiceClassName>
  <Goal>
    <PercentileResponseTime>
      <Importance>Medium</Importance>
      <ResponseTime>00:00:00.200</ResponseTime>
      <Percentile>80</Percentile>
    </PercentileResponseTime>
  </Goal>
</ServicePolicyGoal>
<ServicePolicyGoal>
  <ServiceClassName>SC_Plants_default</ServiceClassName>
  <Goal>
    <AverageResponseTime>
      <Importance>Medium</Importance>
      <ResponseTime>00:00:00.100</ResponseTime>
    </AverageResponseTime>
  </Goal>
</ServicePolicyGoal>
<ServicePolicyGoal>
  <ServiceClassName>SC_Plants_login</ServiceClassName>
  <Goal>
    <PercentileResponseTime>
      <Importance>Medium</Importance>
      <ResponseTime>00:00:00.150</ResponseTime>
      <Percentile>90</Percentile>
    </PercentileResponseTime>
  </Goal>
</ServicePolicyGoal>
<ServicePolicyGoal>
  <ServiceClassName>SC_Trade3_quotes</ServiceClassName>
  <Goal>
    <AverageResponseTime>
      <Importance>Medium</Importance>
      <ResponseTime>00:00:00.800</ResponseTime>
    </AverageResponseTime>
  </Goal>
</ServicePolicyGoal>
<ServicePolicyGoal>
  <ServiceClassName>SC_Trade3_portfolio</ServiceClassName>
  <Goal>
    <AverageResponseTime>
      <Importance>Medium</Importance>
      <ResponseTime>00:00:00.500</ResponseTime>
    </AverageResponseTime>
  </Goal>
</ServicePolicyGoal>
<ServicePolicyGoal>
  <ServiceClassName>SC_Trade3_default</ServiceClassName>
  <Goal>
    <AverageResponseTime>
      <Importance>Medium</Importance>
      <ResponseTime>00:00:01.200</ResponseTime>
    </AverageResponseTime>
  </Goal>
</ServicePolicyGoal>
<ServicePolicyGoal>
  <ServiceClassName>SC_Trade3_profile</ServiceClassName>
  <Goal>
```

```
            <AverageResponseTime>
              <Importance>Medium</Importance>
              <ResponseTime>00:00:01.000</ResponseTime>
            </AverageResponseTime>
          </Goal>
        </ServicePolicyGoal>
        <ServicePolicyGoal>
          <ServiceClassName>SC_Trade3_account</ServiceClassName>
          <Goal>
            <AverageResponseTime>
              <Importance>Medium</Importance>
              <ResponseTime>00:00:00.700</ResponseTime>
            </AverageResponseTime>
          </Goal>
        </ServicePolicyGoal>
        <ServicePolicyGoal>
          <ServiceClassName>SC_Trade3_shopping</ServiceClassName>
          <Goal>
            <PercentileResponseTime>
              <Importance>Medium</Importance>
              <ResponseTime>00:00:10.000</ResponseTime>
              <Percentile>90</Percentile>
            </PercentileResponseTime>
          </Goal>
        </ServicePolicyGoal>
        <ServicePolicyGoal>
          <ServiceClassName>SC_Plants_general</ServiceClassName>
          <Goal>
            <Discretionary/>
          </Goal>
        </ServicePolicyGoal>
        <ServicePolicyGoal>
          <ServiceClassName>SC_Trade3_general</ServiceClassName>
          <Goal>
            <Discretionary/>
          </Goal>
        </ServicePolicyGoal>
        <ServicePolicyGoal>
          <ServiceClassName>EWLM Service Class</ServiceClassName>
          <Goal>
            <Discretionary/>
          </Goal>
        </ServicePolicyGoal>
      </ServicePolicyGoals>
    </ServicePolicy>

  </ServicePolicies>

  <Applications>

    <Application>
      <ApplicationName>IBM Webserving Plugin</ApplicationName>
      <GroupName>(\*)</GroupName>
      <Description>Filters for IHS/Apache, IIS and the IBM WebSphere webserver
plugins.</Description>
    </Application>

    <Application>
      <ApplicationName>(\*)</ApplicationName>
      <GroupName>(\*)</GroupName>
```

```
    <Description>Default Application</Description>
  </Application>

</Applications>

<TransactionClasses>

  <TransactionClass>
    <Name>TC_Trade3_buy</Name>
    <Description>Buy transaction class for Trade3.</Description>
    <ApplicationName>IBM Webserving Plugin</ApplicationName>
    <GroupName>(\*)</GroupName>
    <ServiceClassName>SC_Trade3_profile</ServiceClassName>
    <Filter>
      <FilterType>EWLM:URI</FilterType>
      <FilterOperation>stringMatch</FilterOperation>
      <FilterValue>/trade(\*)</FilterValue>
      <Filter>
        <FilterType>QueryString</FilterType>
        <FilterOperation>stringMatch</FilterOperation>
        <FilterValue>action=buy(\*)</FilterValue>
      </Filter>
    </Filter>
  </TransactionClass>

  <TransactionClass>
    <Name>TC_Trade3_sell</Name>
    <Description>Sell transaction class for Trade3.</Description>
    <ApplicationName>IBM Webserving Plugin</ApplicationName>
    <GroupName>(\*)</GroupName>
    <ServiceClassName>SC_Trade3_shopping</ServiceClassName>
    <Filter>
      <FilterType>EWLM:URI</FilterType>
      <FilterOperation>stringMatch</FilterOperation>
      <FilterValue>/trade(\*)</FilterValue>
      <Filter>
        <FilterType>QueryString</FilterType>
        <FilterOperation>stringMatch</FilterOperation>
        <FilterValue>action=sell(\*)</FilterValue>
      </Filter>
    </Filter>
  </TransactionClass>

  <TransactionClass>
    <Name>TC_Trade3_account</Name>
    <Description>Account transaction class for Trade3.</Description>
    <ApplicationName>IBM Webserving Plugin</ApplicationName>
    <GroupName>(\*)</GroupName>
    <ServiceClassName>SC_Trade3_account</ServiceClassName>
    <Filter>
      <FilterType>EWLM:URI</FilterType>
      <FilterOperation>stringMatch</FilterOperation>
      <FilterValue>/trade(\*)</FilterValue>
      <Filter>
        <FilterType>QueryString</FilterType>
        <FilterOperation>stringMatch</FilterOperation>
        <FilterValue>action=account(\*)</FilterValue>
      </Filter>
    </Filter>
  </TransactionClass>
```

```
<TransactionClass>
  <Name>TC_Trade3_register</Name>
  <Description>Register transaction class for Trade3.</Description>
  <ApplicationName>IBM Webserving Plugin</ApplicationName>
  <GroupName>(\*)</GroupName>
  <ServiceClassName>SC_Trade3_shopping</ServiceClassName>
  <Filter>
    <FilterType>EWLM:URI</FilterType>
    <FilterOperation>stringMatch</FilterOperation>
    <FilterValue>/trade(\*)</FilterValue>
    <Filter>
      <FilterType>QueryString</FilterType>
      <FilterOperation>stringMatch</FilterOperation>
      <FilterValue>Full+Name=(\*)</FilterValue>
    </Filter>
  </Filter>
</TransactionClass>

<TransactionClass>
  <Name>TC_Trade3_update_profile</Name>
  <Description>Update profile transaction class for Trade3.</Description>
  <ApplicationName>IBM Webserving Plugin</ApplicationName>
  <GroupName>(\*)</GroupName>
  <ServiceClassName>SC_Trade3_profile</ServiceClassName>
  <Filter>
    <FilterType>EWLM:URI</FilterType>
    <FilterOperation>stringMatch</FilterOperation>
    <FilterValue>/trade(\*)</FilterValue>
    <Filter>
      <FilterType>QueryString</FilterType>
      <FilterOperation>stringMatch</FilterOperation>
      <FilterValue>userID=(\*)</FilterValue>
    </Filter>
  </Filter>
</TransactionClass>

<TransactionClass>
  <Name>TC_Trade3_home</Name>
  <Description>Home transaction class for Trade3.</Description>
  <ApplicationName>IBM Webserving Plugin</ApplicationName>
  <GroupName>(\*)</GroupName>
  <ServiceClassName>SC_Trade3_default</ServiceClassName>
  <Filter>
    <FilterType>EWLM:URI</FilterType>
    <FilterOperation>stringMatch</FilterOperation>
    <FilterValue>/trade(\*)</FilterValue>
    <Filter>
      <FilterType>QueryString</FilterType>
      <FilterOperation>stringMatch</FilterOperation>
      <FilterValue>action=home(\*)</FilterValue>
    </Filter>
  </Filter>
</TransactionClass>

<TransactionClass>
  <Name>TC_Trade3_logout</Name>
  <Description>Logout transaction class for Trade3.</Description>
  <ApplicationName>IBM Webserving Plugin</ApplicationName>
  <GroupName>(\*)</GroupName>
```

```
  <ServiceClassName>SC_Trade3_default</ServiceClassName>
  <Filter>
    <FilterType>EWLM:URI</FilterType>
    <FilterOperation>stringMatch</FilterOperation>
    <FilterValue>/trade(\*)</FilterValue>
    <Filter>
      <FilterType>QueryString</FilterType>
      <FilterOperation>stringMatch</FilterOperation>
      <FilterValue>action=logout(\*)</FilterValue>
    </Filter>
  </Filter>
</TransactionClass>

<TransactionClass>
  <Name>TC_Trade3_portfolio</Name>
  <Description>Portfolio transaction class for Trade3.</Description>
  <ApplicationName>IBM Webserving Plugin</ApplicationName>
  <GroupName>(\*)</GroupName>
  <ServiceClassName>SC_Trade3_portfolio</ServiceClassName>
  <Filter>
    <FilterType>EWLM:URI</FilterType>
    <FilterOperation>stringMatch</FilterOperation>
    <FilterValue>/trade(\*)</FilterValue>
    <Filter>
      <FilterType>QueryString</FilterType>
      <FilterOperation>stringMatch</FilterOperation>
      <FilterValue>action=portfolio(\*)</FilterValue>
    </Filter>
  </Filter>
</TransactionClass>

<TransactionClass>
  <Name>TC_Trade3_quotes</Name>
  <Description>Quotes transaction class for Trade3.</Description>
  <ApplicationName>IBM Webserving Plugin</ApplicationName>
  <GroupName>(\*)</GroupName>
  <ServiceClassName>SC_Trade3_quotes</ServiceClassName>
  <Filter>
    <FilterType>EWLM:URI</FilterType>
    <FilterOperation>stringMatch</FilterOperation>
    <FilterValue>/trade(\*)</FilterValue>
    <Filter>
      <FilterType>QueryString</FilterType>
      <FilterOperation>stringMatch</FilterOperation>
      <FilterValue>action=quotes(\*)</FilterValue>
    </Filter>
  </Filter>
</TransactionClass>

<TransactionClass>
  <Name>TC_Plants_login</Name>
  <Description>Login transaction class for Plants.</Description>
  <ApplicationName>IBM Webserving Plugin</ApplicationName>
  <GroupName>(\*)</GroupName>
  <ServiceClassName>SC_Plants_login</ServiceClassName>
  <Filter>
    <FilterType>EWLM:URI</FilterType>
    <FilterOperation>stringMatch</FilterOperation>
    <FilterValue>/PlantsByWebSphere(\*)</FilterValue>
    <Filter>
```

```
      <FilterType>QueryString</FilterType>
      <FilterOperation>stringMatch</FilterOperation>
      <FilterValue>action=login(\*)</FilterValue>
    </Filter>
  </Filter>
</TransactionClass>

<TransactionClass>
  <Name>TC_Plants_register</Name>
  <Description>Register transaction class for Plants.</Description>
  <ApplicationName>IBM Webserving Plugin</ApplicationName>
  <GroupName>(\*)</GroupName>
  <ServiceClassName>SC_Plants_login</ServiceClassName>
  <Filter>
    <FilterType>EWLM:URI</FilterType>
    <FilterOperation>stringMatch</FilterOperation>
    <FilterValue>/PlantsByWebSphere(\*)</FilterValue>
    <Filter>
      <FilterType>QueryString</FilterType>
      <FilterOperation>stringMatch</FilterOperation>
      <FilterValue>action=register(\*)</FilterValue>
    </Filter>
  </Filter>
</TransactionClass>

<TransactionClass>
  <Name>TC_Plants_getimage</Name>
  <Description>Get image transaction class for Plants.</Description>
  <ApplicationName>IBM Webserving Plugin</ApplicationName>
  <GroupName>(\*)</GroupName>
  <ServiceClassName>SC_Plants_default</ServiceClassName>
  <Filter>
    <FilterType>EWLM:URI</FilterType>
    <FilterOperation>stringMatch</FilterOperation>
    <FilterValue>/PlantsByWebSphere(\*)</FilterValue>
    <Filter>
      <FilterType>QueryString</FilterType>
      <FilterOperation>stringMatch</FilterOperation>
      <FilterValue>action=getimage(\*)</FilterValue>
    </Filter>
  </Filter>
</TransactionClass>

<TransactionClass>
  <Name>TC_Plants_productdetail</Name>
  <Description>Product detail transaction class for Plants.</Description>
  <ApplicationName>IBM Webserving Plugin</ApplicationName>
  <GroupName>(\*)</GroupName>
  <ServiceClassName>SC_Plants_default</ServiceClassName>
  <Filter>
    <FilterType>EWLM:URI</FilterType>
    <FilterOperation>stringMatch</FilterOperation>
    <FilterValue>/PlantsByWebSphere(\*)</FilterValue>
    <Filter>
      <FilterType>QueryString</FilterType>
      <FilterOperation>stringMatch</FilterOperation>
      <FilterValue>action=productdetail(\*)</FilterValue>
    </Filter>
  </Filter>
</TransactionClass>
```

```
<TransactionClass>
  <Name>TC_Plants_gotocart</Name>
  <Description>Go to cart detail transaction class for Plants.</Description>
  <ApplicationName>IBM Webserving Plugin</ApplicationName>
  <GroupName>(\*)</GroupName>
  <ServiceClassName>SC_Plants_shopping</ServiceClassName>
  <Filter>
    <FilterType>EWLM:URI</FilterType>
    <FilterOperation>stringMatch</FilterOperation>
    <FilterValue>/PlantsByWebSphere(\*)</FilterValue>
    <Filter>
      <FilterType>QueryString</FilterType>
      <FilterOperation>stringMatch</FilterOperation>
      <FilterValue>action=gotocart(\*)</FilterValue>
    </Filter>
  </Filter>
</TransactionClass>

<TransactionClass>
  <Name>TC_Plants_initcheckout</Name>
  <Description>Init checkout transaction class for Plants.</Description>
  <ApplicationName>IBM Webserving Plugin</ApplicationName>
  <GroupName>(\*)</GroupName>
  <ServiceClassName>SC_Plants_shopping</ServiceClassName>
  <Filter>
    <FilterType>EWLM:URI</FilterType>
    <FilterOperation>stringMatch</FilterOperation>
    <FilterValue>/PlantsByWebSphere(\*)</FilterValue>
    <Filter>
      <FilterType>QueryString</FilterType>
      <FilterOperation>stringMatch</FilterOperation>
      <FilterValue>action=initcheckout(\*)</FilterValue>
    </Filter>
  </Filter>
</TransactionClass>

<TransactionClass>
  <Name>TC_Plants_shopping</Name>
  <Description>Shopping transaction class for Plants.</Description>
  <ApplicationName>IBM Webserving Plugin</ApplicationName>
  <GroupName>(\*)</GroupName>
  <ServiceClassName>SC_Plants_shopping</ServiceClassName>
  <Filter>
    <FilterType>EWLM:URI</FilterType>
    <FilterOperation>stringMatch</FilterOperation>
    <FilterValue>/PlantsByWebSphere(\*)</FilterValue>
    <Filter>
      <FilterType>QueryString</FilterType>
      <FilterOperation>stringMatch</FilterOperation>
      <FilterValue>action=shopping(\*)</FilterValue>
    </Filter>
  </Filter>
</TransactionClass>

<TransactionClass>
  <Name>TC_Trade3_general</Name>
  <Description>General transaction class for Trade3.</Description>
  <ApplicationName>IBM Webserving Plugin</ApplicationName>
  <GroupName>(\*)</GroupName>
```

```
    <ServiceClassName>SC_Trade3_general</ServiceClassName>
    <Filter>
      <FilterType>EWLM:URI</FilterType>
      <FilterOperation>stringMatch</FilterOperation>
      <FilterValue>/trade(\*)</FilterValue>
    </Filter>
  </TransactionClass>

  <TransactionClass>
    <Name>TC_Plants_general</Name>
    <Description>General transaction class for Plants.</Description>
    <ApplicationName>IBM Webserving Plugin</ApplicationName>
    <GroupName>(\*)</GroupName>
    <ServiceClassName>SC_Plants_general</ServiceClassName>
    <Filter>
      <FilterType>EWLM:URI</FilterType>
      <FilterOperation>stringMatch</FilterOperation>
      <FilterValue>/PlantsByWebSphere(\*)</FilterValue>
    </Filter>
  </TransactionClass>

  <TransactionClass>
    <Name>default-IBM Webserving Plugin</Name>
    <ApplicationName>IBM Webserving Plugin</ApplicationName>
    <GroupName>(\*)</GroupName>
    <ServiceClassName>EWLM Service Class</ServiceClassName>
    <Filter>
      <FilterType>(\*)</FilterType>
      <FilterOperation>stringMatch</FilterOperation>
      <FilterValue>(\*)</FilterValue>
    </Filter>
  </TransactionClass>

  <TransactionClass>
    <Name>EWLM Default Application Transaction Class</Name>
    <Description>Transaction class for the default application</Description>
    <ApplicationName>(\*)</ApplicationName>
    <GroupName>(\*)</GroupName>
    <ServiceClassName>EWLM Service Class</ServiceClassName>
    <Filter>
      <FilterType>(\*)</FilterType>
      <FilterOperation>stringMatch</FilterOperation>
      <FilterValue>(\*)</FilterValue>
    </Filter>
  </TransactionClass>

</TransactionClasses>

<ServerPlatforms>

  <ServerPlatform>
    <ServerPlatformName>AIX</ServerPlatformName>
    <Description>AIX 5.2.F filters</Description>
  </ServerPlatform>

  <ServerPlatform>
    <ServerPlatformName>Windows</ServerPlatformName>
    <Description>Windows 5.0 (Windows 2000) and 5.2 (Windows 2003) filters</Description>
  </ServerPlatform>
```

```
  </ServerPlatforms>

  <ProcessClasses>

    <ProcessClass>
      <Name>PC_InternetExplorer</Name>
      <Description>Internet Explorer process class for Windows 2003 Server.</Description>
      <ServerPlatformName>Windows</ServerPlatformName>
      <ServiceClassName>SC_InternetExplorer</ServiceClassName>
      <Filter>
        <FilterType>ExecutablePath</FilterType>
        <FilterOperation>stringMatch</FilterOperation>
        <FilterValue>C:\Program Files\Internet Explorer\IEXPLORE.EXE</FilterValue>
      </Filter>
    </ProcessClass>

    <ProcessClass>
      <Name>PC_Calc_batch</Name>
      <Description>Calc batch process class for AIX.</Description>
      <ServerPlatformName>AIX</ServerPlatformName>
      <ServiceClassName>SC_Calc_batch</ServiceClassName>
      <Filter>
        <FilterType>ExecutablePath</FilterType>
        <FilterOperation>stringMatch</FilterOperation>
        <FilterValue>/usr/sbin/calc.batch</FilterValue>
      </Filter>
    </ProcessClass>

  </ProcessClasses>

</DomainPolicy>
```

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

IBM Redbooks

For information about ordering these publications, see “How to get IBM Redbooks” on page 391. Note that some of the documents referenced here might be available in softcopy only.

- *Tuning IBM @server* xSeries Servers for Performance, SG24-5287
- *OS/390 Workload Manager Implementation and Exploitation*, SG24-5326
- *IBM @server* zSeries Connectivity Handbook, SG24-5444
- *Developing Workflows and Automation Packages for IBM Tivoli Intelligent ThinkDynamic Orchestrator*, SG24-6057
- *Grid Services Programming and Application Enablement*, SG24-6100
- *Implementing Systems Management Solutions using IBM Director*, SG24-6188
- *Grid Computing with the IBM Grid Toolbox*, SG24-6332
- *IBM Enterprise Workload Manager*, SG24-6350
- *z/OS Version 1 Release 6 Implementation*, SG24-6377
- *zSeries Application Assist Processor (zAAP) Implementation*, SG24-6386
- *Implementing VMware ESX Server 2.1 with IBM TotalStorage FAStT*, SG24-6434
- *On Demand Operating Environment: Creating Business Flexibility*, SG24-6633
- *On Demand Operating Environment: Managing the Infrastructure (Virtualization Engine Update)*, SG24-6634
- *Effective zSeries Performance Monitoring Using Resource Measurement Facility*, SG24-6645
- *Introduction to Grid Computing with Globus*, SG24-6895
- *Partitioning Implementation for IBM p5 Servers*, SG24-7039
- *Advanced POWER Virtualization on IBM p5 Servers:Introduction and Basic Configuration*, SG24-7940
- *Provisioning On Demand Introducing IBM Tivoli Intelligent ThinkDynamic Orchestrator*, SG24-8888
- *Fundamentals of Grid Computing*, REDP-3613
- *IBM @server* BladeCenter Systems Management with IBM Diector V4.1 and Remote Deployment Manager V4.1, REDP-3776
- *Monitoring Redundant Uninterruptible Power Supplies Using IBM Director*, REDP-3827
- *Server Consolidation with VMware ESX Server*, REDP-3939
- *VMware ESX Server: Scale Up or Scale Out?*, REDP-3953
- *Enterprise Workload Manager - Interpreting Control Center Performance Reports*, REDP-3963

© Copyright IBM Corp. 2005. All rights reserved.

- *Enterprise Workload Manager Overview*, REDP-3989
- *Virtual Partition Manager: a Guide to Planning and Implementation*, REDP-4013
- *Hardening the EWLM Performance Data*, REDP-4018
- *System Programmer's Guide to: Workload Manager*, SG24-6472
- *Virtualization and the On Demand Business*, REDP-9115
- *An Introduction of IBM Tivoli Provisioning Manager Workflows*, TIPS0447
- *Enterprise Workload Manager z/OS Support*, REDP-3993

Other publications

These publications are also relevant as further information sources:

- *IBM @server i5 and iSeries System Handbook:IBM i5/OS Version 5 Release 3, GA19-5486-26*
- *Tivoli Provisioning Manager Installation Guide for UNIX*, GC32-1615
- *TPM Installation Guide for Linux*, GC32-1616
- *WebSphere MQ System Administration Guide*, SC34-6068
- *DB2 Command Reference*, S20H-4645

Online resources

These Web sites and URLs are also relevant as further information sources:

- To get the latest information about the IBM Virtual Machine Manager (VMM)

  http://www-307.ibm.com/pc/support/site.wss/document.do?lndocid=MIGR-56914
- Autonomic computing

  http://www-03.ibm.com/autonomic/index.shtml
- OPAL Orchestration and Provisioning Automation Library (OPAL

  http://www.developer.ibm.com/isv/tivoli/workflow.html
- ARM enabled products

  http://www-1.ibm.com/servers/eserver/about/virtualization/technical/faqs.html#ewlm
- Service Access Point IDI

  http://www-1.ibm.com/solutions/sap/us/en/xslpage/xmlid/26681Description3
- The latest information about the supported operating systems are updated on the Information Center:

  http://publib.boulder.ibm.com/infocenter/eserver/v1r1/en_US/index.htm?info/icmain.htm
- Planning Advisor

  http://publib.boulder.ibm.com/infocenter/eserver/v1r1/en_US/index.htm?info/veinfo/planner.htm
  http://publib.boulder.ibm.com/infocenter/eserver/v1r1/en_US/index.htm?info/icmain.htm
- Virtualization Engine Web Site

  http://www-1.ibm.com/servers/eserver/about/virtualization/getting_started.htmlVEWebsite

- Passport Advantage

  http://www.ibm.com/software/passportadvantage

- Fix Packs

  http://techsupport.services.ibm.com/server/VirtualizationEngine

How to get IBM Redbooks

You can search for, view, or download Redbooks, Redpapers, Hints and Tips, draft publications and Additional materials, as well as order hardcopy Redbooks or CD-ROMs, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

Numerics

10831 Hotfix 194
5665-M76 75
5685-M76 75
5724-i15 75
5724-i71 73, 75
5724-i72 73–75
5733-VE1 25, 74–75
5765-G22 75

A

Advanced POWER Virtualization, *see* APV
aio 122, 151
AIX 13, 22–25, 27, 31, 51, 63–64, 73–76, 88, 97, 102, 118, 120, 122–123, 125, 132, 134, 139, 143–145, 147–152, 155–156, 158–159, 245–246, 248–249, 251, 254–255, 262–263, 265–266, 274–275, 280, 283, 285–286, 288–289, 294–295, 301, 347, 353, 359–360, 362
AIX Toolbox 150
AIX utilities
  SMIT 120, 122–123, 151–152
Apache HTTP Server 39
applying patches 107, 154
APV 22, 24
architecture 11, 354, 356, 362
ARM 13, 37–39, 118, 147, 192, 275–276, 278–279
asynchronous I/O device, *see* aio 151
automation 6–7, 10, 12, 23, 44, 251, 300–301, 317, 319
  availability 7
  optimization 7
  policy-based orchestration 8
  provisioning 8
  security 7
autonomic computing technologies 10–11, 17, 314
  common console 10
  data collection 10
  policy and information model 10
  resource provisioning 10
  solution installation 10
availability 7
average response time 38

B

best practices 349, 354
BladeCenter 21, 31, 44, 143, 226, 262, 291, 306
business
  consequences 6
  priority 5–6, 10
  requirements 9
  services management 8

C

Capacity Manager 31
CBE 10
CIM 14, 186, 267
cluster 28–29, 43, 48, 143, 245, 275, 286, 289
cluster management functions 23
commands
  aio 122–123, 151
  cfgdb 133, 185
  changeCC 142, 190
  chgaio 123, 151
  chmod 144, 353
  configwizard 119, 140–141
  configwizardDM 140, 187
  configwizardMS 144, 191
  create 153
  createDM 140–141, 190
  createMS 144
  DB2licm 194
  db2sampl 133
  disable 156
  displayCC 142–143, 145, 190
  displayDM 141, 363–364
  EWLMAixMS 144
  EWLMWinMS 191
  export 43, 96, 154, 216, 363–364
  genevent 318
  geninstall 120
  grep 135
  gunzip 236
  ibmdirctl 125, 154–155
  installp 123, 134, 151
  installVEAix 139
  installvewin 180, 183, 186
  ldapcfg 105, 107, 124
  ldapsearch 127
  ldapxcfg 95, 105, 152
  lslpp 143–144, 154
  mkitab 153
  ps 135, 285
  pwdadm 152
  reboot 153, 174, 182, 184, 237, 369
  reinit 247
  rpm 94, 151, 237
  setupaix 155
  startdm 144–145, 190
  startdmservice 190
  startMS 145
  startserver 175
  startWAS 190
  stopServer 155
  tar 134, 155, 220, 236
  unzip 155, 219
  versioninfo 182
  wpversioninfo 182

© Copyright IBM Corp. 2005. All rights reserved.

xmlimport 43, 158, 199, 247, 294
xterm 120
Common Base Event, *see* CBE
common console format 10
Common Information Model, *see* CIM
common runtime 10, 58, 60–61, 74–75, 81, 109, 263, 351
configuring
database 106
DB2 database 195
DCM 112
EWLM Domain Manager 140
EWLM Managed Servers 144
LDAP directory 125
prerequisites 100, 150
Tivoli Directory Server 94, 105, 122, 151, 195
Tivoli Provisioning Manager 197, 294
Virtualization Engine console 125, 137
Control Center 36–37, 51
Coupling Facility 28
Covalent Enterprise Ready Server 39
creating workflows 251
credential 303, 305, 358
CSM 22–23, 51, 53, 130, 137
CUoD 25, 27
customizing workflows 251
Cygwin 159, 164, 192, 197, 292, 302

D

dashboard 12, 15, 55
data center resources 294
database
configuring 106
DB2 39, 51, 59–60, 62, 74, 87, 92, 95, 97, 99, 102–106, 109, 119, 123, 132–133, 147, 149, 151–154, 160, 164, 170, 174, 185, 193–195, 245–246, 252, 263–266, 279–281
DB2 requirements 97
DCM 41, 43, 109, 112–113, 150, 158, 164, 197, 199, 248–249, 292, 294–295, 299
defining a threshold 324
DeviceID 256
discretionary 38
Distributed Management Task Force, Inc, *see* DMTF
DLPAR 23, 25–27
partitioning 22
DMTF 14
Domain Manager, *see* EWLM
domain policy 37
dynamic logical partitioning, *see* DLPAR

E

encryption 93, 184, 220–221, 354–355, 359–360
Enterprise Workload Manager, *see* EWLM
ESX, *see* VMware
Ethernet 22, 24–25, 27, 228, 367
Event Action Plan Builder 317, 327–331, 333, 337–338, 342
Event Action Plans 45
EWLM 13, 15–16, 19, 25, 36, 38, 40, 48, 63–65, 73–74, 81, 118–119, 134, 139–141, 143–146, 166, 168, 186–187, 189–190, 202, 204, 212, 217–219, 222, 244, 246, 250–254, 256–258, 263–266, 268, 279–284, 286–289, 291, 306, 349, 353, 355–357, 360, 362–365, 370, 373, 377
Control Center 36–37, 51, 119, 141–143, 145–147, 190, 222, 252, 265–266, 268, 280–283, 286–289, 360, 362
Domain Manager 36–37, 62, 74, 118–119, 139–147, 164, 187, 189–191, 213, 217, 221–222, 246, 251, 257, 262–263, 266, 287, 289, 353, 355–356, 360–364
Managed Server 37–38, 62, 76, 119, 139, 143–144, 164, 191, 204, 213, 222, 252–253, 257, 274, 281, 289
Managed Servers 144
Management Domain 36, 222, 236, 263, 266, 274, 280, 286, 289, 367
planning the installation 139
policy 36, 246, 253, 266, 327, 349, 353–355, 360, 377
security settings 188
starting 190
topology 36

F

fault tolerance 166
Fibre Channel 22, 24–25
fix download number 77
fix pack 76–77, 109–110, 158

G

GDPS 27
graphical installer 87, 104, 108, 119–120, 123, 151–152, 155, 368
Grid Toolbox 15, 61, 74–75, 81
52P1214 75
52P1263 75
Guest LANs 27

H

Hardware Management Console, *see* HMC
hardware requirements, Virtualization Engine 91
Health Center 49, 52, 96
dashboard 55
Resource Health 53–54
Watch List 55, 57–58
HiperSockets 12, 27–28
HMC 22
host partition 22
HTTP Server 39
HTTP server 262
Hypervisor 15, 23, 25, 223, 312

I

i5/OS 22–23, 25, 27, 31, 64, 73–74, 76
IBM Director 15–17, 22, 25, 27, 31, 33, 53, 74–75, 81, 86–88, 90–91, 93–94, 96–97, 99, 110, 119, 130, 133–137, 166, 168, 180–181, 183–185, 202, 204, 212,

216, 218–221, 226–240, 262–264, 291–293, 297, 299–300, 302–304, 306–310, 312, 314, 317–319, 322, 325–327, 329, 332–338, 341, 343, 347–349, 353–360, 367, 369, 373
Agent 45, 53, 88–89, 97–98, 132, 134, 164, 186, 204, 213, 219–221, 227, 229, 233–234, 236–238, 263–264, 310, 312, 314, 319, 321, 335, 341, 353, 359, 369
Console 31, 45, 51, 97, 118, 134–136, 219, 221, 226, 229, 231–235, 237–238, 311–313, 315–320, 322, 324–326, 328–329, 332, 336–339, 341, 367, 373
installing 96, 132
installing the code 98
Multiplatform 15, 18, 25, 44–45, 51, 53, 55, 58, 60, 63–65, 74–75, 97, 110, 202, 263
security 98
Server 45, 62, 96, 118, 132–133, 137, 181, 186, 204, 219, 221, 227, 229–231, 233–234, 293, 302, 314, 318, 323, 325–327, 332, 334, 353, 355, 358
IBM Director Execution Window 315
IBM Directory Server 62, 89, 91
installing 95
requirements 97
IBM HTTP Server 39
IBM Software Maintenance 74
Information Center 2–3, 68, 81, 83, 88, 91, 93–96, 110, 120–121, 123, 125–127, 133, 139, 161, 164, 171, 192, 345–349, 360
installing
10831 Hotfix 194
applying patches 107
DB2 133
DB2 license 194
EWLM Domain Manager 139, 187
EWLM Managed Server 143, 191
fix pack 109
IBM Director 96, 98
IBM Director Agent 132, 134
IBM Director Server 132–133
IBM Directory Server 95
Linux 101
on AIX 120
planning 2, 122, 127, 139
post-installation tasks 158
pre-installation recommendations 120
pre-installation recommendations, BladeCenter with Linux 87
prerequisites 86, 100–101, 118, 150, 153, 192
process 119
process for BladeCenter with Linux 87
requirements for BladeCenter on Linux 91
testing the installation 199
Tivoli Directory Server 94, 122–123, 151, 194
Tivoli Provisioning Manager 99, 147, 149, 155, 192, 197
verification 127
verifying 131, 135, 144
Virtualization Engine console 95, 121, 127
Virtualization Engine console bridge 96, 137
Virtualization Engine console code 95
Virtualization Engine console on BladeCenter with Linux 93
VMware VirtualCenter 371
WebSphere fix pack 194
InstallShield 108, 240
integrated shared infrastructure 31
Integrated Solutions Console, *see* ISC
Intel 12–13, 31, 101, 108, 118, 225
Intelligent Resource Director, *see* IRD
Intelligent ThinkDynamic Orchestrator 102
interview 71
IRD 27, 29
ISC 10, 47–48, 51, 59, 62, 73–74, 87, 90, 131, 164, 168, 180–181, 213–216, 218, 269, 350, 357–358
iSeries 12–13, 21, 25–27, 39, 44, 51, 58, 64, 74, 101, 108, 201–203, 205–208, 211, 221, 349

J

Java 30–31, 39, 44

L

Launchpad 49–50
LDAP 48, 58, 87, 89–90, 92, 94, 99, 107, 121–122, 124–129, 137, 152–154, 156–157, 160, 164, 168–172, 174–176, 180, 193, 203–205, 207–209, 213–216, 218, 263, 269, 354, 356, 358, 360
Lightweight Third Party Authentication, *see* LTPA
Linux 22–29, 31–32, 51, 65, 75, 88–89, 96–97, 101, 108–110, 112, 118, 150, 213, 220–221, 227–228, 236, 246, 291, 301, 311, 321–322, 347, 360, 369
load balancer 42
Logical Partition Validation Tool (LVT) 23, 26
logical partitioning, *see* LPAR
logical topology 78
logical unit numbers 25
Lotus Domino Enterprise Server 39
LPAR 12, 22–23, 26–29, 120
partitioning 12, 23–24, 148, 223, 262–264
LTPA 93, 96, 137, 164, 168, 180, 214–216, 354–355, 357–358
LTPA key information 130
LUN 25

M

managed systems 220, 294
Management Central for iSeries 53
Management Domain, *see* EWLM
management source 48, 86, 96, 130, 137–138, 269–270, 273, 357–358
management sources 48
MAPE loop 11
Micro-Partitioning 22–23, 25–26
Microsoft Internet Information Services (IIS) 39
Microsoft Virtual Server 33, 225–226, 229–230, 310–311, 313
Multiple Device Manager 51, 53
MVS 28

N

NetWare 32, 110, 294
network requirements, Virtualization Engine 93
NIC 248–249, 295, 297–299

O

OASIS 14
On Demand Operating Environment 6–7, 10, 12, 14, 17
OPAL 9, 251
open standard 6, 12, 15, 18
openSSH 150, 158–159, 192
OptiConnect, *see* Virtual OptiConnect
optimization 7, 18–19, 279
Orchestration and Provisioning Automation Library, *see* OPAL
Organization for the Advancement of Structured Information Standards, *see* OASIS
OS/400 27, 64, 201–207, 211, 213, 218, 220–222, 347

P

P2V, *see* VMware
Parallel Sysplex 27–29
partition 12, 24, 26–27, 100, 376
partition host 22
Partition Load Manager, *see* PLM
partitioning
  DLPAR 22
Passport Advantage 73, 76, 87, 211, 350
patches 107
  applying 154
percentile response time 38
performance 6, 28, 30, 228, 244–245, 253, 258, 264–267, 275, 279, 281–282, 286–287, 289, 349, 355, 368
performance goal, *see also* WLM and EWLM 38, 253, 265–266, 286
Performance Index, *see* PI
Performance Monitoring Infrastructure, *see PMI*
physical topology 79
PI 38, 253–254, 282, 285–286, 289
PICS 193–194
planning advisor 67–68, 71, 76, 78–81, 88, 91, 93, 96, 203–204, 211, 348, 350
  code 73
  Deploy tab 89
  Install tab 90
  installation process 80
  Passport Advantage 73
  tasks 80
platforms class 37
PLM 13, 22–24
plug-in 52, 245, 275–279, 300
PMI 275–278
policy 246, 253, 266, 349, 353–355, 360, 377
  domain 37
policy-based orchestration 8
pool
  logical 12, 32
portlet 49, 52, 219, 274
post-installation tasks 158
POWER Hypervisor 23
POWER4 23, 26
POWER4+ 23, 26
POWER5 13, 22–27, 143, 148, 262–264
PR/SM 27, 29
pre-installation recommendations 120
prerequisites 100, 153, 192
priority, see also business 288–289, 319
process class 37
provisioning 8, 16, 40, 147, 153, 263, 266, 275, 285–287, 289, 291–294, 296–302, 306–307
pSeries 12–13, 21–22, 24, 44, 63, 75, 143, 147, 349

R

RDM 291–294, 297, 299–301, 304, 306–307
RDS, *see* Resource Dependency Service
Red Hat 101–103, 109, 112
reporting 28
requirements
  database 97
  prerequisites 100
requirements, IBM Directory Server 97
resiliency 5, 7, 13
Resource Dependency Service 16, 62
Resource Health 53–54
Resource Manager 44
resource monitoring 45
resource pools 42
RSA 170, 248–249, 303, 305, 307, 362, 364

S

SAN Virtualization Controller 25
SAP 248, 303, 305
SCSI 24–25, 27
SCSI adapters 22
security 6–7, 79, 129, 141, 170, 180, 188, 214–216, 347, 350, 353–355, 357–360, 362
  IBM Director 98
security requirements, Virtualization Engine 93
sensor 237
servers 42
Service Access Points 249
service class 37, 253, 258, 280, 284, 286–287, 377
service definition, *see also* WLM and EWLM 28
service level 6, 266, 288
service-oriented architecture, *see* SOA
shared infrastructure 31
SMIT 120, 122–123, 151–152
SNMP 101, 193, 248
SOA 6, 15
SOAP 44
software requirements
  Virtualization Engine 92
SoftwareID 256
SQL 89, 132–133, 185, 370
SSL 36, 93, 98, 100, 126, 141, 144, 170, 174, 205, 208, 237, 354–355, 358–365
Start Your Engines (SYE) workshops 2

starting
EWLM 190
Tivoli Directory Server 196
subnet 42
Sun ONE Web Server, Enterprise Edition 39
Sun Solaris 74
SuSE 24, 64, 86, 88, 97, 99, 102–103, 109–110
switch 42
Systems Provisioning Management Server 62

T

tasks
planning advisor 80
tc driver 44
testing the installation 199
threshold 324
tioadmin 159–160
tioadmin user shell 159
Tivoli Change and Configuration Management 62
Tivoli Directory Server 59–60, 73–74, 79, 87, 93, 95, 102–105, 119, 125, 148–149, 151, 153–155, 164, 168–176, 179–180, 182, 185, 193–197, 246
configuring 94, 122
installing 94, 122–123
Tivoli Directory Server Rebidding Tool 123
Tivoli Intelligent Orchestrator 9, 87, 154, 248
Tivoli Monitoring 48, 53
Tivoli NetView 101, 157
Tivoli Provisioning Manager 9, 15–16, 40, 42–44, 58, 63–65, 74, 76, 86–87, 89–90, 93, 99–100, 108, 112, 148–151, 153–156, 158–160, 164, 166, 192, 194, 197, 199, 244, 262, 274, 294, 304
Automation Package 44, 301
DCM 41, 109, 112–113, 150, 158, 164, 197, 199, 248–249, 292, 294–295, 299
deployment engine 44, 251, 305
installing 99, 147
workflow 43–44, 199, 223, 250–252, 255–257, 274, 289, 293–294, 299–300, 302–303, 306–308, 349
Tivoli SAN Manager 51
Tivoli Software Installer 102, 153
Tivoli Storage Resource Manager 51
tool
planning advisor 67–68
toolkit 170
tools 202, 247, 283, 292, 301, 367, 372
topology 118, 148, 153–154, 163, 203, 223, 226, 228–230, 235, 245–246, 253–254, 258, 262–263, 266, 280, 284, 286, 288, 292, 312–313, 348
EWLM 36
logical 78
physical 79
TotalStorage Enterprise Storage Server 25
TPF 28
Trade3 245, 253, 263–266, 275, 280, 282, 284–286, 377
transaction class 37, 280
Transaction Processing Facility, *see* TPF
transitions in a workflow 43

U

Updatewizard, see wizard 155

V

VE Farm, *see* Virtualization Engine Farm
velocity 38
verifying the installation 127, 131, 135, 144
VIO 15, 22, 24–25, 27, 295
VLAN 292, 294–295, 297
VIO Server 24
virtual I/O, *see* VIO
Virtual LANs 12
Virtual Machine Manager 31, 33, 225–235, 238–240, 309–313, 315–316, 338–339, 367
Virtual Machine Resource Manager 27
Virtual OptiConnect 27
Virtual Opticonnect 27
Virtual Partition Manager, *see* VPM
VirtualCenter, *see* VMware
virtualization 12–13, 16, 18, 25, 27, 225–226, 230
evolution of 15
Virtualization Engine 13, 16, 40, 44, 73–74, 76
5665-M76 75
5685-M76 75
5724-i15 75
5724-i71 73, 75
5724-i72 73–75
5733-VE1 74–75
5765-G22 75
common runtime 10, 58, 60–61, 74–75, 81, 109, 263, 351
hardware requirements 91
installation process 80, 83, 90, 119
installation tasks 80
management collection 16, 67
network requirements 93
planning advisor 67–68, 71, 73, 76, 78–81, 88–91, 93, 96, 203–204, 211, 348, 350
pre-installation recommendations 87
requirements 91
security requirements 93
software requirements 92
Systems Edition 16, 21
topology 67, 86
virtual access 15–16
Virtualization Engine Console 262
Virtualization Engine console 10, 25, 46–47, 74, 90, 93, 110, 121, 166, 204–205, 212, 216, 218, 360
bridge 15, 48, 53, 89, 93, 96, 130, 137, 164, 166, 168, 180, 184, 212, 216, 218, 263, 269, 354–357, 359–360
worksheet 96
browser 48
components 47
configuring the LDAP directory 125
Health Center 49, 52, 96
dashboard 55
Monitors 55
Resource health 53–54
Watch List 55, 57–58

installing 93, 95, 121, 127
Launchpad 49–50
LDAP settings 95
managed systems 48
management sources 48
worksheet 94
Virtualization Engine Farm 237, 311, 313, 375–376
Virtualization Engine Installer 74, 139
Virtualization Engine installer 120, 127
Virtualization Engine Management Servers, *see* 5724-i71
VLAN 15, 30, 42, 247–249, 292, 294–295, 297
VLAN, *see* VIO
VMM, *see* Virtual Machine Manager
VMotion, *see* VMware
VMware 13, 31, 33, 225–230, 233–238, 309, 313–314, 317, 319, 327, 340, 343, 367–376
ESX 375
ESX Server 32, 313, 317, 319
GSX Server 32
P2V 32
VirtualCenter 32, 225–230, 233–235, 237, 309–313, 319, 334, 338–340, 367, 370–376
VMotion 32, 314, 317, 376
VMware ESX 225–229, 236–238, 309–312, 314, 327, 334–335, 338, 343, 367–370, 373, 375
VMware workstation 228
Workstation 32
VPM 13, 26
VSE/ESA 29

W

Watch List 55, 57–58
Web Services 96, 339, 358, 372–373
Web services 13–14
Web Services Distributed Management, *see* WSDM
Web-based System Manager 51
WebSphere 48
WebSphere Application Server 30, 39, 48, 59, 62, 73–74, 87, 90, 92, 95, 99–100, 104, 107–108, 110–111, 118, 124–125, 137, 141, 145, 149–151, 153–157, 164, 182, 187, 197, 206–207, 213–214, 216–218, 222, 244–246, 249–252, 254, 256–257, 259, 262–263, 275, 279, 350, 358
WebSphere Portal 47, 131
WebSphere Studio Workload Simulator 244, 246, 253, 264
Windows 31–32, 39, 64, 73–74, 88, 96–97, 110, 118, 132–134, 144, 168–169, 171, 173–175, 182, 185, 190, 202, 211, 219, 221–222, 226–228, 230–234, 238–240, 245–246, 251, 262–265, 267, 280, 285, 291–293, 297, 299–302, 304, 306–308, 310, 312, 314, 321–322, 340–341, 347, 353, 359, 362, 370–375
wizard 140, 142, 144, 155, 164, 166, 168, 179–180, 182–183, 185–186, 190–191, 204–206, 211–213, 217–219, 221–222, 240, 278, 293, 348, 350–351, 375–376
workflow 8, 43–44, 251, 255, 257, 274, 289, 293–294, 299–300, 302–303, 306–308
adding in the Resource Pool 251
allocate servers to clusters 43
assigning initialization 250
best practices 9
configure network devices 43
configure servers 43
creating 251
customized 9, 251
definition of Service Access Point 250
dummy_install 251
dummy_uninstalll 251
engine 223
executing 249
initialization 252
install software and patches 43
installation 252
prebuilt 9
remove servers from clusters 43
run commands 43
software installation 256
status 256
tab 251
testing the installation 199
to start WebSphere Application Server 252
transitions 43
uninstall software 43
worksheet 94, 96, 121
workshops 2
WS Resource Framework specifications, *see* WS-RF
WSDM 14
WS-RF 14

X

xLinux 91, 132
XML 37, 42–43, 113, 199, 247, 277–279, 294, 340, 358, 372–373
xSeries 13, 21, 31, 33, 44, 65, 75, 96–97, 99, 110, 226, 291, 306, 349

Z

z/OS 13, 28–30, 74, 245–246, 263–266, 268, 279
z/VE 28
z/VM 13, 27–29
zAAP 27, 30
zSeries 12–13, 21, 27–30, 63

IBM Virtualization Engine, Version 1: Planning and Installation Guide

(0.5” spine)
0.475”<->0.873”
250 <-> 459 pages

IBM Virtualization Engine, Version 1: Planning and Installation Guide

How to use the Planning Advisor and the Information Center

Installation examples for IBM @server platforms

How to diagnose problems using the IBM Virtualization Engine

The core of the On Demand operating environment strategy is a simple IT infrastructure. *Virtualization* — the ability to separate the direct dependency of an application to a physical resource — is a key fundamental in simplifying the heterogeneous environment in today's businesses. Through virtualization, the enterprise offers:

- A single consolidated view to all available resources regardless of the location.
- Efficient access to manage those resources and to reduce operations and systems management costs.
- Dynamic response to application needs.
- A competitive advantage when gathering and accessing information across the organization.

The IBM Virtualization Engine addresses the virtualization aspects of the On Demand infrastructure.

This IBM Redbook shows you how to:

- Plan and prepare the installation of the IBM Virtualization Engine Version 1.
- Install the IBM Virtualization Engine Version 1 for each type of IBM @server.
- Use the IBM Virtualization Engine Version 1 to address customer pain points.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-6771-00 ISBN 0738492779